

Financial Literacy, Shocks, and Portfolio Adjustments

Sebastian Gomez-Cardona*

*Vancouver School of Economics
University of British Columbia*

[CLICK HERE FOR THE LATEST VERSION](#)

October 20, 2023

Abstract

I develop a life-cycle model of household portfolio decisions that accounts for heterogeneity in financial literacy and employ it to examine portfolio adjustments following household-level shocks. I use variation in unplanned births to parameterize the model and identify the margins of portfolio adjustments following household-level fertility shocks. Empirical evidence suggests that households increase the liquidity of their portfolios following such shocks. Using the model, I compare how households with different financial-literacy levels respond to similar shocks, and I show that higher financial literacy is associated with smoother portfolio adjustments following shock onset. All else equal, the more financially literate households appear less susceptible to the detrimental effects of liquidity constraints and the impact of portfolio-adjustment costs. The interaction between liquidity constraints and financial literacy plays a key role in the model, as it explains the differential speed and direction of portfolio adjustments observed in the data. Counterfactual exercises show that financial literacy mitigates the negative welfare effects of unexpected fertility shocks by at least 20%.

Keywords: Financial literacy, Shocks, Portfolios, Liquidity, Fertility

*Email: sebgomez@student.ubc.ca, sebastiangomez87@gmail.com. I thank Michael Devereux, Giovanni Gallipoli, Viktoria Hnatkovska and Kevin Milligan for their advice and support. I also thank Annamaria Lusardi for various comments and suggestions. I am grateful to participants at the 8th Cherry Blossom Financial Education Institute, the 57th Annual Conference of the Canadian Economics Association, the G53Network Boot Camp on Financial Literacy and Personal Finance, and workshops at the Vancouver School of Economics and University of Victoria.

1 Introduction

Studying the factors that affect household responses to shocks is instrumental to understanding aggregate outcomes and designing effective public policies. To the extent that liquidity needs dictate households' ability to accommodate expenditures, portfolio adjustments are a useful margin to learn about heterogeneity across households. In this paper, I examine the composition of household portfolios and use portfolio adjustments to gauge the importance of financial literacy for consumption and welfare.

An extensive literature examines household responses to income and expenditure shocks, but less evidence is available on the role played by financial literacy in the adjustment process that follows these shocks. Much of the literature on financial literacy focuses on empirical analysis of its correlation to other household characteristics, often using transitory income shocks (e.g., Angrisani et al., 2020). This is partly due to the limited availability of measures of financial competence. This paper attempts to overcome this challenge by combining two datasets, documenting several stylized facts, and using them to parameterize a model of household behavior. I contribute to the fast-growing literature on financial literacy by describing how it relates to portfolio adjustments and by focusing on the mechanisms that lead to such adjustments. Furthermore, the focus on permanent shocks to consumption provides a stark characterization of the benefits of financial literacy following a shock, particularly for poorer households.

I employ changes in fertility across and within households to quantify the relationship between financial literacy and portfolio adjustments after shocks. Using variation in fertility for this purpose provides three advantages. First, it allows the differentiation between planned and unplanned births. I use planned births as a benchmark to parameterize the relationship between fertility and portfolio adjustments in my model. I interpret unplanned births as shocks with two important characteristics: They are exogenous, and they imply large and persistent increases in expenditure. The significance of fertility in household spending makes the analysis of these shocks even more relevant for policymakers. Moreover, the insights gained from analyzing this type of shock can be applied to other contexts, such as health shocks.

The second advantage of using variation in fertility in this way is that it addresses housing, one of the primary assets of household wealth, making it relevant for welfare and aggregate outcomes. The third advantage is that it allows the analysis of how financial literacy can help overcome adjustment costs related to illiquid assets, in this case, real estate. Moreover, I document the adjustment toward housing across households spanning the entire wealth distribution, providing a valuable lens to investigate the dynamic interaction between borrowing constraints and financial literacy.

In the first part of the paper, I use data from the Panel Study of Income Dynamics

(PSID) and the Survey of Consumer Finances (SCF) to document three stylized facts that serve as inputs for parameterizing the model. Both datasets include information on fertility and household assets. Additionally, the PSID has information on fertility shocks and the SCF on financial literacy, so I use both sources to document the empirical facts.

This study leverages the PSID’s capacity to capture housing and other assets and to track households’ portfolios before and after childbirth. I employ questions from the Childbirth and Adoption History module to classify births as planned or unplanned. By comparing households with unplanned births to those without children, I document the first stylized fact: a substantial decrease in the share of housing following a fertility shock. This result could be explained by households’ need to increase the liquidity of their portfolio, given the increase in expenditures caused by an unplanned birth. The absence of information on financial literacy in the PSID calls for a structural model to establish the relationship between financial literacy and portfolio adjustments in response to shocks.

The second stylized fact, using information from the PSID, documents an increase in the share of housing around the birth of children. This implies that planned children, in contrast to unplanned ones, are related to increased housing share within household portfolios. Three interesting details arise about this adjustment. First, it happens for all households across the wealth distribution with varying magnitudes. Those in the middle tercile of the wealth distribution experience the largest increase in housing share within their portfolios. This disparity reflects the varying capacity of these groups to adjust their portfolios. Those at the lower end of the spectrum, near their borrowing limit, have limited flexibility to adjust, whereas those at the upper end are closer to their preferred portfolio composition, requiring less adjustment following childbirth. Parents in the middle of the wealth distribution find themselves in between, necessitating more significant changes with a fertility shift. Second, this adjustment occurs through both the intensive and extensive margins. Households either acquire larger homes or transition from renting to owning. Third, this portfolio adjustment occurs predominantly before childbirth, indicating household financial planning.

The third stylized fact is related to financial literacy and the speed of portfolio adjustments. The SCF includes the “Big 3” financial literacy questions that measure financial literacy.¹ I leverage variation in the age of children to determine to what extent financial competence relates to the speed with which households adjust their portfolios towards housing as children age. I compare financially literate and illiterate parents within wealth terciles.

Financial competence exerts minimal influence over the pace and magnitude of housing-allocation changes in the highest tercile. This pattern could be explained by households

¹Lusardi and Mitchell (2007) developed these questions to measure the understanding of basic financial concepts: compound interest rates, inflation, and diversification. Section 2.1 presents a descriptive analysis, and the questions are available in Appendix A.1.

in this group compensating financial-literacy shortfalls by using financial advisers, allowing them to achieve their preferred portfolio allocation. In the middle of the wealth distribution, financially literate parents are better able to increase their share of housing earlier in their children’s lives. Conversely, at the lower end of the wealth spectrum, financial literacy delays the adjustment toward housing. An interaction between liquidity constraints, adjustment costs in housing, and financial education can explain these results. As long as parents are not liquidity-constrained, they increase their share of housing. However, if constrained, financial literacy plays a role in overcoming the constraints and the housing adjustment costs.

In the second part of the paper, I develop a parsimonious life-cycle model to examine the relationship between financial literacy and portfolio adjustments after a fertility shock. I also leverage the model to analyze the welfare implications of the interaction between fertility and financial literacy. The model incorporates uninsurable income risk, heterogeneity over both fertility and financial literacy, and liquidity constraints and includes two consumption goods: nondurable and housing services. These services can come from owning or renting a house. The model includes two available assets: a financial asset and a durable good (housing) with adjustment costs. Fertility and financial literacy are exogenous and fixed from time 0, although they are drawn from a multivariate distribution that accounts for their empirical correlation. Financial literacy has two levels: low and high. Those with higher financial literacy have access to a higher return on the financial asset.

This parsimonious model replicates the patterns documented in the stylized facts while allowing for the distinction of how parents with different financial-literacy levels adjust their portfolios to fertility shocks. The mechanism involves the interaction between borrowing constraints and the greater ease with which parents possessing higher financial literacy can manage the associated adjustment costs. This interaction replicates the heterogeneity in the speed of adjustment across the wealth distribution, and it also creates variation in how parents face a fertility shock.

When a shock occurs, parents with higher financial literacy can more easily adjust their portfolios due to their higher return on the liquid asset, helping them to cope with the persistent and large increase in expenditures. Financial literacy can partially offset the welfare losses of fertility shocks in the bottom and middle of the wealth distribution. Financial literacy compensates for at least 20% of the welfare losses of households with unplanned children, taking parents of planned children as a benchmark. This result highlights the importance of analyzing the role of financial literacy in protecting households from shocks. As such, financial literacy can be an additional policy tool to improve households’ welfare during their life cycle (Urban et al., 2020) and when facing large shocks, such as fertility or health shocks.

Related Literature. This paper relates to two strands of the financial-literacy literature. The first studies the relationship between financial literacy and shocks. I contribute to this

literature by analyzing a large permanent shock using a structural model. As mentioned before, most of the papers are empirical contributions using survey answers to transitory shocks (e.g., Angrisani et al., 2020; Lusardi, Hasler, and Yakoboski, 2021; Babiarz and Robb, 2014), finding evidence that financially literate individuals are less likely to use high-cost borrowing during shocks (Lusardi and Bassa Scheresberg, 2013) and cope better with aggregate macroeconomic shocks (Klapper, Lusardi, and Panos, 2013).

The second strand of the financial-literacy literature studies portfolio allocation and the speed of portfolio adjustment. I contribute to this literature by including nonfinancial assets in portfolios. Large literatures document the relationship between financial literacy and portfolio allocation (Gaudecker, 2015; Hastings and Mitchell, 2020), the accumulation of wealth throughout the life cycle (e.g., Van Rooij, Lusardi, and Alessie, 2012; Lusardi and Mitchell, 2007; Lusardi, Michaud, and Mitchell, 2017), access to various financial assets such as stocks (Gaudecker, 2015), and increased returns from specific asset classes (Fagereng et al., 2020). On the speed of portfolio adjustments, Bianchi (2018) documents a positive relationship between financial literacy and the frequency of portfolio rebalances on financial portfolios.

The paper is also preceded by a large literature that analyzes the importance of portfolio liquidity (e.g., Kaplan and Violante, 2014; Luetticke, 2021). I add to this literature by including financial literacy in my model as the ability to access a liquid asset with higher return. I also contribute to the literature that analyzes the households' response to shocks (Liu, 2016; Baker et al., 2023; De Ferra, Mitman, and Romei, 2020; Coile and Milligan, 2009) and its implication in macro aggregates (Kaplan and Violante, 2018).

Finally, the paper contributes to the literature on family composition and housing demand during the life cycle. My model builds on the housing demand models of Fernández-Villaverde and Krueger (2011); Bajari et al. (2013) but includes fertility and financial literacy as key characteristics for housing accumulation in the life cycle. On the relationship between family composition and housing, Bacher (2021) documents the importance of marital status on the demand for housing. While I abstract from this channel, my results complement this literature using childbirth to identify fertility shocks and its relationship to housing demand. More specifically, this paper also relates to the literature on fertility and portfolio choice (e.g., Love, 2009; Bogan, 2013; Görlitz and Tamm, 2015), to which I contribute by including nonfinancial assets while they focus on financial portfolios.

The rest of the paper is organized as follows. Section 2 presents the data and summary statistics. Section 3 explains the empirical strategy used to analyze fertility and portfolio choice. Section 4 presents the stylized empirical facts. Section 5 presents the model and its calibration, while Section 6 analyses welfare and policy implications. Finally, Section 7 concludes.

2 Data

I use the PSID and the SCF to document the stylized facts that will serve to calibrate the model. The PSID is an ideal dataset, including information on fertility and household portfolios. It is a long panel, allowing me to follow households several periods before and after children are born. The PSID is conducted every two years, and the information on portfolios is available since 1999. It includes financial assets (stocks, retirement accounts, and deposits—savings and checking accounts, bonds, and mutual funds) and real assets (main residence, other real estate, and private businesses). For each of these assets, I use the net equity owned by the household. So, for example, for the main residence, I take the value of the house minus any outstanding balance of mortgages. On the fertility side, PSID includes information on the number of children living at home, date of birth, and, since 2013, questions on the intention to have a child. This last set of questions allows me to classify births as planned and unplanned, as explained in Section 2.2.1.

Since the PSID does not include questions on financial literacy, I use the SCF. The 2016 and 2019 waves include three questions, the Big 3 developed by Lusardi and Mitchell (2007) and designed to measure the understanding of basic financial concepts. These well-established questions in the literature allow me to classify households between financially literate and illiterate or high and low financial-literacy levels in the model. The SCF provides cross sections, making it impossible to observe the same household before and after the birth of a child. I leverage variation in children’s ages to analyze the path of household’s assets within their portfolios as children grow. The SCF does not include questions that allow me to classify births as planned or unplanned, so estimating the response to fertility shocks by parents with different financial-literacy levels is impossible.

2.1 Financial literacy

The measurement of financial literacy comes from the SCF, using the Big 3. These simple multiple-choice questions measure the knowledge of basic financial concepts related to compound interest rates, inflation, and diversification (Appendix A.1). They were developed by Lusardi and Mitchell (2007) and are widely used in the literature to measure financial literacy. Households are classified into high and low financial literacy depending on whether they answered the three questions correctly. Despite their simplicity, only about 40% of households can answer the three questions correctly. The question with the highest share of correct answers is interest rate (80%), and the one with the least is diversification (62%). Financial literacy positively correlates with wealth (Table 1). While only one in four households in the lowest tercile have a head who is financially literate, at the top this proportion is two in three. It is important to note that the second tercile also presents low levels of financial

Table 1: Financial literacy level

Share answering correctly			
All	Compound interest rate	Inflation	Diversification
0.43	0.80	0.75	0.62
Share answering three questions correctly			
All	Wealth tercile 1	Wealth tercile 2	Wealth tercile 3
0.43	0.26	0.38	0.64

The first panel indicates the share of the population that answers the questions correctly. The first column indicates the portion answering all three questions correctly. The second panel shows the share of households who answer the three questions correctly by wealth tercile. Source: SCF

literacy, with only one in three answering the three questions correctly. This highlights the vulnerability to shocks of those at the bottom and middle of the wealth distribution: the low savings are accompanied by a low understanding of financial concepts, which might prevent them from responding properly to a shock.

2.2 Fertility

One of the advantages of choosing fertility as a magnifying glass to uncover the relationship between financial literacy and portfolio allocation when facing shocks is its prevalence along the wealth distribution. Figure 1a shows the negative relationship between households living with children and wealth, as has been widely documented (e.g., Bar et al., 2018). While around 45% of households at the bottom have a child living with them, this number decreases to 40% at the top. When breaking down the share by the level of financial literacy, it is clear that the decreasing trend is driven mostly by those with low financial literacy. The gap of nearly 10% for poorer households disappears for richer ones. Although fertility is not an endogenous choice in the model, it is modeled to match the correlations observed in the data on wealth and financial literacy.

2.2.1 Unplanned births

Unplanned births can be an important shock for households, as they imply persistently higher expenditures and could eventually push households toward the borrowing constraint. For births after 2011, the PSID includes questions on whether the child was “wanted” or

not by the mother and the father (Appendix A.2 presents the exact questions). I classify births where the mother answered the child was “not wanted” as fertility shocks. This could either mean they did not plan to have children in their lives or at that specific moment. This definition allows me to identify 550 unplanned births. The proportion of unplanned births in this period is close to 20%, similar to the one found by Miller (2011) using the NLSY. Her definition is whether the mother was using contraceptives when she became pregnant. This question is also available in the PSID; with it, I identify 75 unplanned births (Appendix A.3). Given the small sample size, I use the first question and confirm that my results are robust to this alternative definition.

This type of shock is more prevalent among poorer households (also explored in Su and Addo, 2018). While one of every three births is classified as unplanned at the bottom of the wealth distribution, the proportion decreases to one of every 20 at the top (Figure 1b). Therefore, studying how financial literacy can help households, particularly poorer ones, face these shocks through portfolio adjustments could significantly increase their welfare.

2.3 Portfolios

When analyzing household portfolios, it is important to include all types of assets to understand the adjustments around a shock. Following the PSID, I classify assets into safe financial assets, stocks, main residence, other real estate, retirement accounts, and private businesses. In all cases, the asset refers to the net equity position of the household, net of any debt associated with that specific asset. Liquid debt, which is not associated with any asset, is not included in the analysis, as the paper’s interest is on how households allocate available funds into different assets, not the source of such funds. The descriptive and empirical sections focus on two particular assets: safe financial assets and main residence since the adjustment around an unplanned birth is mostly around these two assets. Appendix B.1




Figure 1: Fertility and unplanned births by wealth level
Sources: SCF and PSID

shows the results for the other assets. Although stocks also have a significant change, less than 10% of the households own stocks, and they are mostly concentrated at the top of the wealth distribution.

As mentioned above, housing refers to the net equity value owned by the household in their main residence by subtracting any outstanding mortgage debt from the house value. The safe financial asset includes checking and savings accounts, money market funds, certificates of deposit, government bonds, and treasury bills. This composition comes from the PSID and is the most detailed level of disaggregation. To make results comparable, the definition in the SCF is all those financial assets different from stocks and retirement accounts.

Using information from the SCF, Table 2 shows the share in the portfolio for these two assets. While virtually all households own a safe financial asset, there is heterogeneity in the allocated share. At the bottom of the wealth distribution, this asset is more than 80% of portfolios; for other households, it is around 20%. On the contrary, the share of homeowners varies greatly by wealth, from 13% at the bottom to 93% at the top. The share allocated to this asset is very similar for the two first terciles, with the average household holding around two-thirds in housing. These numbers highlight the importance of accounting for the extensive and intensive margin in the structural model. When buying a house, those at the bottom of the wealth distribution will probably adjust on the extensive margin, while others adjust on the intensive margin, buying a larger house.

Table 2: Portfolios

Share owning the asset				
Asset	All	Wealth tercile 1	Wealth tercile 2	Wealth tercile 3
Safe financial asset	0.99	0.99	0.98	1.00
Housing	0.63	0.13	0.64	0.93
Portfolio share conditional on owning				
Asset	All	Wealth tercile 1	Wealth tercile 2	Wealth tercile 3
Safe financial asset	0.37	0.84	0.24	0.19
Housing	0.41	0.67	0.67	0.29

Summary statistics for the two assets with the largest adjustment to fertility shocks: safe financial asset and housing. The first panel shows the share of households who own each type of asset, by wealth tercile and the aggregate. The second panel shows the average portfolio share of each asset for those households with nonzero values in their portfolios. Source: SCF

Table 3: Correlation between financial literacy (fin. lit.), fertility, and wealth

	Fin. lit.	Fertility	Wealth tercile
Fin. lit.	1.00		
Fertility	-0.04	1.00	
Wealth tercile	0.39	-0.03	1.00

Correlation between financial literacy, fertility, and wealth. All variables are categorical. Financial literacy has two levels, low and high. Fertility is a dummy for having at least a child. Wealth is classified by terciles. Source: SCF

2.3.1 Financial literacy, wealth, and fertility

Three variables play a key role in the model: financial literacy, wealth, and fertility. Table 3 shows the correlation among them. As can be expected and has been extensively documented in the literature, wealth and financial literacy are highly correlated. In the SCF, the correlation is close to 0.4. As shown in Figure 1b, wealth and fertility (including both planned and unplanned births) are negatively correlated, which implies a similar correlation between fertility and financial literacy. These correlations are around -0.03 in the data.

When introducing these three levels of heterogeneity in the analysis of portfolios (Table 4), some salient features emerge, which will be documented robustly in Section 4. First, when comparing financial literacy within the same wealth tercile, those financially literate tend to hold a lower share of their portfolio in housing. Second, controlling for financial literacy and wealth, those households with children tend to have a higher share of their portfolio in housing, and the difference is higher for those with greater financial literacy.

3 Empirical strategy

The empirical section presents my analysis of how portfolios adjust around births. The adjustment can come from intensive or extensive margins. In the former, the variable of interest is the share of asset k in the portfolio. In contrast, for the latter, the variable of interest is a dummy denoting whether a household owns asset k . Let $y_{i,t}^{*k}$ denote the latent variable of asset k in household i during year t . This latent variable can take any real value, while the observed variables are functions of this latent variable.

Let $y_{i,t}^{T,k}$ denote the share of asset k and $y_{i,t}^{L,k}$ a dummy indicating whether household i owns asset k . These variables are defined as functions of $y_{i,t}^{*k}$:

Table 4: Share of housing

Wealth	Tercile 1	Tercile 2	Tercile 3
Low fin. lit.	0.07	0.44	0.37
No children	0.06	0.43	0.36
Children	0.08	0.46	0.38
High fin. lit.	0.08	0.39	0.23
No children	0.06	0.36	0.22
Children	0.13	0.44	0.24

Average share of housing by fertility, financial literacy and wealth. The average includes households with 0 on their portfolio. Source: SCF

$$y_{i,t}^{T,k} = \begin{cases} 0 & y_{i,t}^{*k} \leq 0 \\ y_{i,t}^{*k} & 0 \leq y_{i,t}^{*k} \leq 1 \\ 1 & y_{i,t}^{*k} \geq 1 \end{cases} \quad (1)$$

$$y_{i,t}^{L,k} = \begin{cases} 0 & y_{i,t}^{*k} \leq 0 \\ 1 & y_{i,t}^{*k} > 0 \end{cases} \quad (2)$$

To analyze the intensive and extensive margins using $y_{i,t}^{T,k}$, I estimate a random-effects Tobit model for the PSID or a standard Tobit model when using the SCF since it presents cross sections. To analyze only the extensive margin, I estimate logits. The empirical section uses three types of regressions:

1. Portfolio adjustments around births
2. Portfolio differences between parents living with children and other households
3. Differences in the adjustment speed between financially literate and illiterate parents.

Although all regressions follow the same logic, the empirical specifications and the control groups differ slightly. All regressions include ten controls in the term $X_{i,t}$: age, age squared, years of education, gender, marital status, state, number of adults in the household, income tercile, wealth tercile, and year fixed effects. Including age and its square in the regression adjusts for asset k share life-cycle trajectories. Estimating the effect by wealth group would include the interaction between the variable of interest and the wealth tercile. In that case, the control group would be those in the same wealth tercile.

3.1 Unplanned births

Control groups are defined in two ways. A household with an unplanned birth can be compared to those with at least one birth or to those who have not had a child and are “old enough” to assume they do not plan to. In the first case, the estimated regression is in equation (3).

$$y_{i,t}^{*k} = \beta^k X_{i,t} + \sum_{j=-N_0}^{N_1} \gamma_{0,j}^k \mathbb{1}\{\text{birth}_{i,t+j}\} + \sum_{j=b, a} \gamma_j^k \mathbb{1}\{\text{birthunplanned}_{i,j}\} + \epsilon_{i,t}^k \quad (3)$$

The first term holds the control variables, the second controls for the path around having a child, while the last term estimates the effect three years before (b) and after (a) having a fertility shock. More than the coefficients γ_b^k and γ_a^k , the interest is on their difference, which captures whether the shock affected asset k . When using as a control group those who never had a child, the regression simplifies to equation (4).

$$y_{i,t}^{*k} = \beta^k X_{i,t} + \sum_{j=b, a} \gamma_j^k \mathbb{1}\{\text{birthunplanned}_{i,j}\} + \epsilon_{i,t}^k \quad (4)$$

Since there is no need to account for the birth of children other than those unplanned, in this case, the control group is all the observations of households whose head is observed at least once being at least 40 years or older. It is assumed that by this moment if they have not had children, they do not plan to. So, the control group is composed of households whose portfolios were never and will never be influenced by fertility changes. Of the observed births, only 10% correspond to a head of the household who is 40 or older at a child’s birth.

3.2 Children living at home

When estimating the effect of a “static” variable ($z_{i,t}$), such as having a child living at home, the regression takes the form of equation (5). The parameter of interest is γ^k , which denotes the effect of having a child at home on asset k . In this regression, the control group is those households without a child living at home.

$$y_{i,t}^{*k} = \beta^k X_{i,t} + \gamma^k z_{i,t} + \epsilon_{i,t}^k \quad (5)$$

3.3 Birth of first child

The other type of effect estimated is around a child’s birth. I focus on the first child since this likely involves the largest portfolio adjustments. In such a case, the specification is that of equation (6). The controls are the same as described above, and the coefficients of interest are γ_j^k . These coefficients represent the trajectory of asset k around the birth of a child

with respect to the control group. Using all available information in the PSID since 1999, the effects estimated here include all births, *planned* and *unplanned*, as the distinction is only available since 2011. I set $N_0 = 4$ and $N_1 = 2$, corresponding to eight and four years before and after the birth. The choice of N_1 is determined to make it comparable to a similar estimation of unplanned births for which there are not enough observations to make $N_1 = 3$.

$$y_{i,t}^{*k} = \beta^k X_{i,t} + \sum_{j=-N_0}^{N_1} \gamma_j^k \mathbb{1}\{\text{birth}_{i,t+j}\} + \epsilon_{i,t}^k \quad (6)$$

Defining the control group involves two options. First, use the complete sample so the control group is households without a child. The disadvantage in this case is the presence of some households who might be planning to have a child in the future and are adjusting their portfolio in advance. This would create a downward bias in the estimates. The second option is to define the control as those who have not had a child and whose head is “too old” to do so, just as in the unplanned case explained in Section 3.1. The advantage is that using the observations of these households when they were younger, the control group does not include any portfolio adjustment anticipating a birth. The main text presents the results for the whole sample as the control group, but the results are robust to using the second control group.

3.4 Portfolio-adjustment speed and financial literacy

Financial literacy is only observed in the SCF, which offers cross sections. The adjustment speed is analyzed after a child’s birth and within wealth terciles. Children’s young lives are split into four periods (groups): ages 0–4, 5–9, 10–14, and 15–17. The estimated regression, equation (7), includes the interaction of these variables (j) with the financial-literacy level (f) and the wealth tercile (l). In this case, the coefficients of interest are $(\gamma_{l,f,j}^k)$, which denote the effect of having a child on asset k with respect to households without children in the same wealth and financial-literacy groups.

$$y_i^{*k} = \beta^k X_i + \sum_{l=1}^3 \sum_{f=1}^2 \sum_{j=1}^4 \gamma_{l,f,j}^k \mathbb{1}\{\text{child group} = j\} \mathbb{1}\{\text{finlit group} = f\} \mathbb{1}\{\text{wealth tercile} = l\} + \epsilon_i^k \quad (7)$$

4 Stylized facts on household portfolios

This section presents the three stylized facts related to fertility and portfolio adjustments. The first is about fertility shocks: unplanned births. I document a decline in the share of housing, both in the intensive and extensive margins and for all wealth levels, especially the lowest one. The second is related to planned and unplanned births pooled together. I

document an increase in housing by families with children, which, given the first fact, is mostly driven by families with planned births. They have a planning horizon of around two years, which shows that planned births are expected, and households adjust their portfolios a couple of years in advance. The third stylized fact is related to financial literacy and portfolio-adjustment speed. Using variation in the age of children, I document heterogeneity in the relationship between speed of adjustment and financial literacy.

I use the second and third facts to calibrate the structural model in Section 5. The results focus on housing and, in some cases, the safe financial asset, as these are the assets with the largest adjustments when facing a fertility shock (see Appendix B.1). I also perform the robustness checks presented in Appendix B. Figures and tables present the marginal effect, calculated with the command *margins* in Stata.

4.1 Fertility shocks

Fertility shocks imply unexpected permanent increases in household expenditures. The empirical results support the hypothesis that the need for liquidity increases as the housing share decreases while the safe financial asset's share increases. Figure 2 compares the share of housing between those with fertility shocks and households without children, as explained in Section 3.1. Data comes from the PSID, and the periods “before” and “after” are three observations before and after the shock. Since the PSID is biannual, this is equivalent to six years before and after. Appendix B.3 shows the path year by year, in which it is clear that the decrease in the share of housing comes right after the birth of the unplanned child, but it continues a few years later. Therefore, households adjust immediately but continue to do so for several years after the shock. This is an important characteristic of fertility shocks, as they are permanent and can lead to large and permanent adjustments, giving more room for financial literacy to play a role.

Figure 2 shows how the share of housing decreases after the shock while the share of the safe financial asset increases. The shares of both assets were not significantly different from those of the control group before the shock. However, after the shock, housing decreased by 6%, and the financial asset increased around equivalently. A similar result is obtained when the control group consists of households with planned births (Appendix B.4).

4.1.1 Extensive margin

The previous results come from Tobit, which include adjustments in the intensive and extensive margins. Figure 3 comes from a logit, in which the dependent variable is a dummy for owning a house. The effect is significant at 10% and shows an extensive margin adjustment after a shock. Households with an unplanned child are 3% less likely to own a house six years




Figure 2: Share before and after a fertility shock. Control group: households without children. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.




Figure 3: Owning a house before and after a fertility shock. Control group: households without children. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.

after the shock. This points to the importance of including rental markets in the structural model, as this is a channel through which households can adjust their portfolios to increase liquidity.

4.1.2 Effect by wealth tercile

When decomposing the effect by wealth tercile—that is, comparing households with a fertility shock to those without children in the same wealth tercile—the direction is the same, and the




Figure 4: Share before and after a fertility shock by wealth tercile. Control group: households without children. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.

effect is significant for most cases. Figure 4 shows that the increase in the safe financial asset is close to being significant for the first two terciles. This fact is consistent with households' increase in the need for liquidity, creating an adjustment toward liquid assets.

The next two subsections document stylized facts around births, both planned and unplanned, that will calibrate the structural model and shed light on the relationship between births and portfolio adjustments.

4.2 Fertility and housing

The results in this subsection come from the PSID and are related to the difference in portfolios between households living with children and those without and the path around the birth of a child. The first is intended to uncover the relationship between fertility and portfolios when including nonfinancial assets, where I document a preference for having a larger share in housing. The second is related to the adjustment around the birth of a child. I document an adjustment towards housing that begins before the first child's birth, on average between two and four years before. Since this includes both planned and unplanned births, the results imply that for planned births, households increase the share of housing in advance of their birth. This contrasts with unplanned births presented in Section 4.1, which decrease the share of housing around the birth.

Table 5: Child living at home and portfolio shares

	Share		Own
	Housing	Safe fin. asset	Housing
Child at home	0.069*** (0.002)	-0.029*** (0.003)	0.031*** (0.004)
Controls	X	X	X
<i>N</i>	59,532	60,192	73,616

Control group: households without children. Effect on dependent variable. Calculated using *margins* in Stata. Standard errors clustered by household in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.

4.2.1 Family size

The first two columns in Table 5 show the adjustment towards housing and away from the liquid financial asset for households with children living with them. This is a family-size effect, as households tend to allocate a larger share of their portfolio to housing when they have more members. Other factors can affect the demand for housing and its share in portfolios. For example, Bacher (2021) documents marital status as an important driver of this adjustment. Therefore, I do not claim this is the largest housing determinant, but an empirical relationship I account for in my model.

4.2.1.1 Extensive margin

The third column in Table 5 documents that the extensive margin plays an important role in increasing the share of housing in portfolios. This means that not only do households with children have larger houses, they are also more likely to own a house compared to those who do not have children living with them. The effect on the extensive margin is almost half the total effect, as households with children are 3% more likely to own a house.

4.2.1.2 By wealth tercile

Figure 5 shows the differences across the wealth distribution on the adjustment. In the middle of the wealth distribution, the adjustment is the largest, while at the bottom is the lowest. This is consistent with poorer households being constrained and less likely to own a house,




Figure 5: Housing share by wealth tercile. Control group: households without children living at home. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.

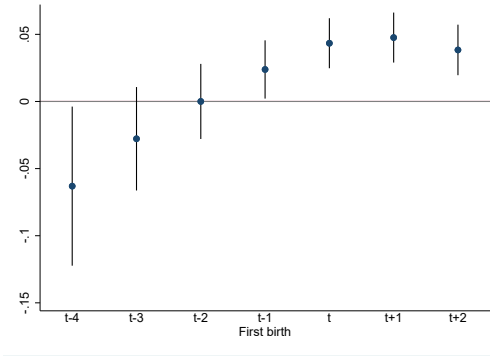
even if they want to, while those at the top do not adjust as much as they probably had large houses even in the absence of children.

4.2.2 Planning horizon


This fact documents how the adjustment occurs around the birth of children. The results include planned and unplanned births, so the former dominate the effect, as they are around four-fifths of the sample. I use the birth of the first child as this is the one with the largest adjustment. Figure 6 shows the adjustments around the first child's birth. In contrast to the fertility shocks, households with planned children increase the share of housing in their portfolio while decreasing the safe financial assets. The increase in housing begins a couple of years before childbirth (every period is equivalent to two years since the PSID is biannual).

4.2.2.1 Extensive margin

Figure 7a shows that in the case of the first planned child, the adjustment in the extensive margin is more important than the aggregate effect. While the adjustment of having a child living at home comes around 50% from buying a house, in the case of the first child, it is close to 80% (4% in Figure 7a out of 5% in Figure 6). This implies that adjustments after the first child are more on the intensive margin. This is another argument favoring the inclusion of a rental market in the structural model.




(a) Housing




(b) Safe financial asset

Figure 6: Share around birth of first child. Control group: households without children. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.



(a) Own a house



(b) Share of housing

Figure 7: Birth of first child. Control group: all households. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.

4.2.2.2 By wealth tercile

The adjustment towards housing occurs at all wealth levels, as shown in Figure 7b. This proves that the decrease in housing when facing a fertility shock is not driven by the poorest households, who are more likely to face these shocks. Their magnitude is the largest in terms of the increase in the share of housing.


4.3 Financial literacy and adjustment speed

The speed of adjustment with which households can make portfolio adjustments can play a crucial role when facing shocks. To calibrate the model, I use all births, both planned and unplanned, to uncover the relationship between financial literacy and the speed of portfolio adjustments. Recall that it is impossible to differentiate in the SCF by type of birth. The results in this section are about the share of assets in portfolios. Appendix B.7 shows that part of the differences in adjustment speed comes from the extensive margin.


Figures 8, 9, and 10 show the fertility differences within wealth terciles and financial-literacy levels. The analysis is throughout children’s lives (Section 3.4). Financial literacy is related to adjustment speed in the first two terciles, those at greatest risk when facing shocks. In the bottom tercile (Figure 8), those with low financial literacy invest more in housing at childrens’ early ages but then lag. In contrast, high-financial-literacy parents increase the share of deposits during the children’s first years and then, when they are 10, increase the share of housing. Investing in deposits allows them to have a much higher share in housing




Figure 8: Share of asset in portfolio. Lowest tercile of the wealth distribution. Control group: households without children with the same financial literacy level and same wealth tercile. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.




(a) Safe fin. asset. Low financial literacy



(b) Safe fin. asset. High financial literacy



(c) Housing. Low financial literacy



(d) Housing. High financial literacy

Figure 9: Share of asset in portfolio. Middle tercile of the wealth distribution. Control group: households without children with the same financial literacy level and same wealth tercile. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.

later in children’s lives. The difference in the patterns for the two levels of financial literacy is explained in the model by those with higher financial literacy having a higher return on financial assets, which allows them to overcome the transaction costs to buy a house.

This mechanism can also explain the pattern in the middle third of the wealth distribution (Figure 9). Those with high financial literacy can own a house earlier in their children’s lives, so their share of housing increases earlier in their children’s lives. The difference between the first and second tercile is the speed at which they can save enough to overcome transaction costs. While those at the bottom are closer to the borrowing constraint, so the saving room is smaller, those in the middle of the wealth distribution could have accumulated enough resources to increase the share of housing earlier in their child’s life.

At the top of the wealth distribution (Figure 10), there is no difference in portfolio shares between those with different levels of financial literacy. In this part of the distribution, financial literacy does not play a role in overcoming the transaction costs: Parents adjust their portfolios when having a child at the same speed.

Results are easier to understand when taking the difference between the effect for those

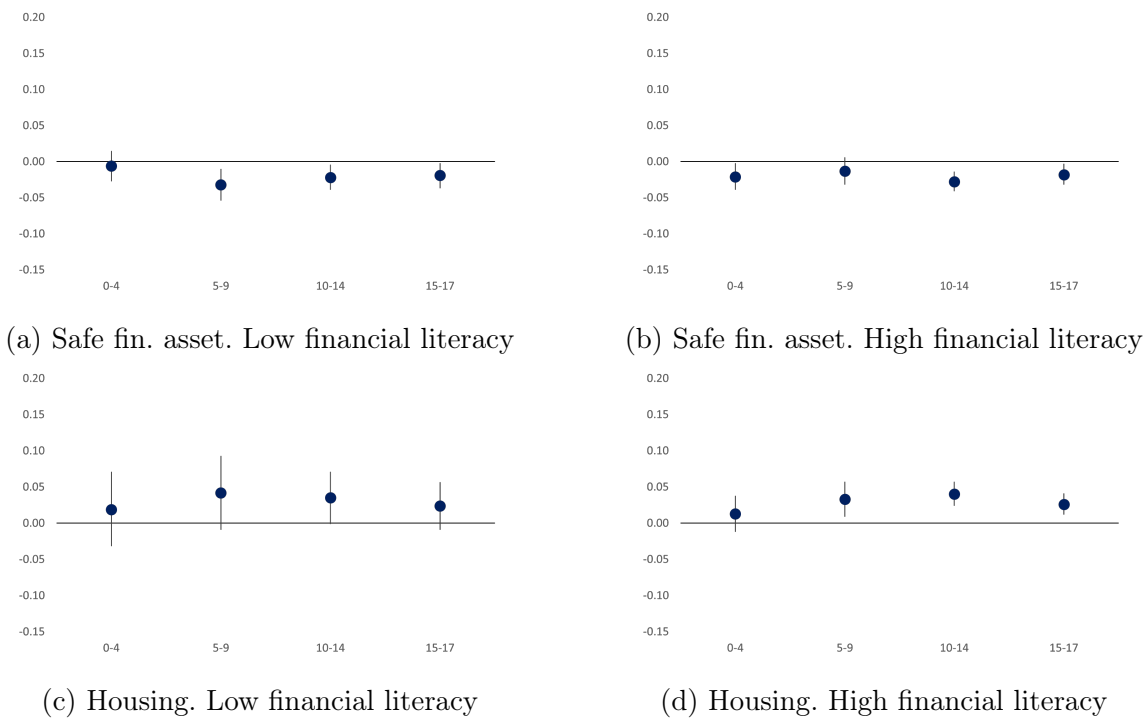


Figure 10: Share of asset in portfolio. Top tercile of the wealth distribution. Control group: households without children with the same financial literacy level and same wealth tercile. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.

with high and low financial literacy. Figure 11 shows how those at the bottom of the wealth distribution with high financial literacy hold less housing in their portfolio when their children are young but increase their share later in their lives. So, financial literacy is related to a delay in the adjustment towards housing for poorer households. In the middle of the wealth distribution, those with high financial literacy hold a higher share of housing earlier in their children’s lives, so financial literacy accelerates the adjustment toward housing. Finally, at the top, there is no difference between parents with high and low financial literacy.

5 Model

This section presents the structural life-cycle model that captures the patterns described above and helps understand the mechanism through which financial literacy affects portfolio adjustments after shocks. It is also used to assess the relative importance of financial literacy when comparing households with different fertility, and in general, could be used to expand on the design of policies that want to capture the dynamics of portfolio adjustment to

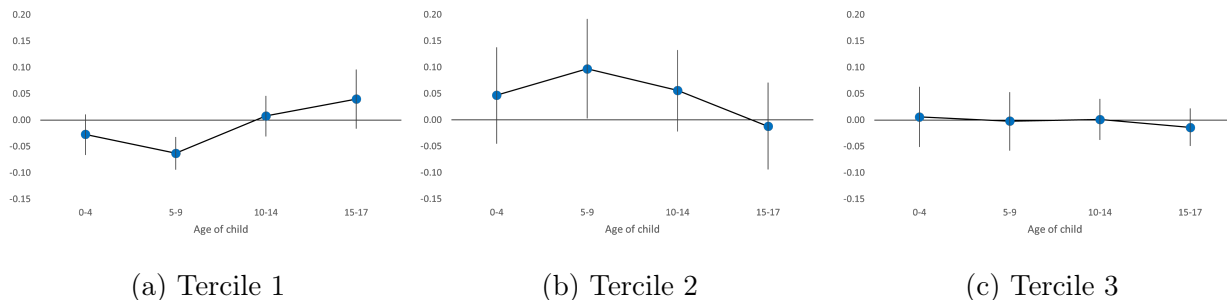


Figure 11: Share of housing in portfolio. Double difference by fertility and financial literacy by wealth tercile. Calculated using *margins* in Stata. Controls head of household: age, age squared, years of education, gender, marital status. Controls household: state, number of adults, year fixed effects, income tercile, and wealth tercile

different shocks for households with different levels of financial literacy. My model is one of the first structural models including financial literacy, thus exploring the mechanism through which financial literacy affects household finances. It moves forward from Lusardi, Michaud, and Mitchell (2017) in two important ways: i) the addition of nonfinancial assets and ii) household composition with the inclusion of fertility. In this regard, it is also one of the first on how financial literacy can affect the choice of illiquid assets, such as housing, which is paramount to analyzing shocks and their implications. In this sense, the model builds on a large literature of housing in life-cycle models (e.g., Fernández-Villaverde and Krueger, 2011: who study the choice between liquid and illiquid assets).

Although it is not part of the core of this paper, the model also permits the analysis of the relationship between fertility and portfolios and its interaction with financial literacy. Love (2009) analyzes the interaction between children and financial portfolios: How children can tilt portfolios between safe and risky assets. Although my model is simplified from the fertility side, it augments the analysis of illiquid assets and their interaction with fertility.

5.1 Model setup

In the life-cycle, partial-equilibrium model, households are composed of adults (a) and, potentially, children (d). Every period in the model is two years, and adults live for 23 periods, ages 20 to 66. Households consume nondurable goods (c) and housing services (m). The housing services can be acquired by renting ($s = 0$) or owning ($s = 1$) a house. The size of the rented (l) or owned (h) house can take a specific set of discrete values, as is standard in the literature (e.g., Bajari et al., 2013). The durable good, housing, is subject to adjustment costs, which is also very common in the literature (e.g., Bajari et al., 2013; Fernández-Villaverde and Krueger, 2011; Fella, 2014).

Households can use either the durable good to save or a risk-free financial asset (b). The return on this asset (R^f) depends on the household's financial-literacy level (f), which can be either low or high. Although in a risky asset, this is the same approach used by Lusardi, Michaud, and Mitchell (2017), supported by growing empirical evidence. Fagereng et al. (2020), in Table 6, find that a degree in economics or business affects the returns of different assets. Bianchi (2018) documents how investors with higher financial literacy have higher returns on their investments, even after adjusting for risk. The difference can be up to 40 basis points between the highest and lowest level of financial literacy. Deuffhard, Georgarakos, and Inderst (2019) also document an increase of 12% in the savings return when financial literacy increases by one standard deviation.

In the model, fertility is exogenous. That is, it is not a choice of the households, and it is deterministic. There are three types of households in terms of fertility. First, those who have planned children. Since the first period, these households know that they are having one child and the period in which they will have it. Those who do not have children also know this since the first period, and they plan accordingly. Finally, households with unplanned children think they will never have a child, but in a given period, they have one unexpectedly. Planned and unplanned births can only occur in the fourth and seventh periods. These are equivalent to ages 26 and 32, the first and second terciles of the parent's age. Children live for nine periods, 18 years, with their parents.

Although financial literacy and fertility are exogenous in the model, I acknowledge their correlation and draw both from a copula with the same correlation between the two variables as in the data. The copula also includes a third characteristic, the initial permanent shock to income, explained below. This means these three characteristics are exogenous but follow the pattern observed in the data. Despite financial literacy being exogenous, endogenizing it would only reinforce the channel. Gallipoli and Gomez-Cardona (2023) show that children do not affect the accumulation of financial literacy, and the process is related to learning by doing. In this model's context, those with higher initial financial literacy would likely invest more and thus accumulate more financial literacy.

Finally, note that the model does not include mortgages, so there is no possibility of using housing as collateral. This simplification is based on three facts. First, the interest of the model is in the use of resources, not so much on the source. Second, the empirical counterpart of housing is the value of the house net of mortgage balances. So, in the model context, h should be interpreted not as the house's value but as the proportion which is equity for the household. Finally, a third fact is the mechanism. In this model, the mechanism is the difference in the financial asset's return. If mortgages were included, the natural assumption would be introducing a differential in the mortgage rate between those with high and low financial literacy. The mechanism would be similar, as those with higher financial literacy

would more easily overcome the costs of buying a house as their lending rate would be lower.

5.2 Household problem

Households enter every period with savings in the financial asset (b_t), the size of the owned house (h_{t-1}), which can be zero if they rented in $t-1$, a level of income y_t , and whether they have a child living at home ($d_t = 0.25$)² or not ($d_t = 0$). In the first stage, the household chooses to rent ($s = 0$) or own ($s = 1$) to consume housing services.

$$V_t(b_t, h_{t-1}, y_t, d_t, f) = \max_{s_t \in \{0,1\}} (1 - s_t)V_t^r(b_t, h_{t-1}, y_t, d_t) + s_t V_t^o(b_t, h_{t-1}, y_t, d_t) \quad (8)$$

If the household decides to rent ($s_t = 0$), they choose the size of savings in the financial asset (b_{t+1}), and the consumption of the nondurable good and housing services for adults and children (c_t^a, c_t^d, l_t^a , and l_t^d). The utility function $u(c^a, c^d, h^a, h^d, d)$ is a composite function based on Fernández-Villaverde and Krueger (2011); Attanasio et al. (2016). The outer one is a CRRA that controls the intertemporal substitution, where $1/\sigma$ is the elasticity of intertemporal substitution. The inner function is a CES with elasticity $1/(1 - \rho)$ among durable and nondurable goods. The parameters α_1 and ψ_1 govern the relative weight of the consumption of adults and children (if any) within a household, while the parameters α_2 and ψ_2 govern the weights of housing services. $\phi(h_{t-1}, h_t)$ is the adjustment cost, which depends on the size of the previously owned house and the size of this period's owned house. In this case, $h_t = 0$ as the household chose to rent. This function can capture many transaction costs in the housing market, such as commissions. p_r denotes the relative price of renting compared to owning, which is normalized to one.

$$\begin{aligned} V_t(b_t, h_{t-1}, y_t, d_t, f) = & \max_{c_t^a, c_t^d, b_{t+1}, l_t^a, l_t^d} \frac{(\alpha_1 (c_t^a)^\rho + \psi_1 d_t (c_t^d)^\rho + \alpha_2 (g_f m_t^a)^\rho + \psi_2 d_t (g_f m_t^d)^\rho)^{(1-\sigma)/\rho}}{1 - \sigma} \\ & + \beta \mathbb{E}_t[V_{t+1}(b_{t+1}, 0, y_{t+1}, d_{t+1}, f)] \\ \text{s.t.} & \\ & b_{t+1} + p_r l_t^a + p_r d_t l_t^d + c_t^a + d_t c_t^d + \phi(h_{t-1}, 0) = R_b^f b_t + h_{t-1} + y_t \\ & m_t^a = l_t^a \\ & m_t^d = d_t l_t^d \\ & b_{t+1} \geq 0 \\ & l_t^a \geq 0 \\ & l_t^d \geq 0 \end{aligned}$$

²A child is usually taken as half an adult for consumption purposes (e.g., Attanasio et al., 2016). Here, each household consists of a couple, so a child would be equivalent to 0.25 of a couple.

In the case of owning ($s_t = 1$), the problem is very similar, although housing services do not depend on the size of the house owned. There is a linear relationship composed of a constant ($taste_o(h_t, d_t)$) and a multiplicative term ($1 + g(d_t)$), which denotes how much more households with children enjoy housing compared to renting. This includes services such as the school's quality, renovating without the landlord's approval, and avoiding eviction.

$$\begin{aligned}
V_t(b_t, h_{t-1}, y_t, d_t, f) = & \max_{c_t^a, c_t^d, b_{t+1}, h_t^a, h_t^d} \frac{(\alpha_1(c_t^a)^\rho + \psi_1 d_t (c_t^d)^\rho + \alpha_2 (g_f m_t^a)^\rho + \psi_2 d_t (g_f m_t^d)^\rho)^{(1-\sigma)/\rho}}{1 - \sigma} \\
& + \beta \mathbb{E}_t[V_{t+1}(b_{t+1}, h_t, y_{t+1}, d_{t+1}, f)] \\
\text{s.t.} & \\
b_{t+1} + h_t^a + d_t h_t^d + c_t^a + d_t c_t^d + \phi(h_{t-1}, h_t) = & R_b^f b_t + h_{t-1} + y_t \\
m_t^a = taste_o(h_t, d_t) + (1 + g(d_t))h_t^a & \\
m_t^d = (1 + g(d_t))h_t^d & \\
h_t = h_t^a + d_t h_t^d & \\
m_t = m_t^a + d_t m_t^d & \\
m_t = taste_o(h_t, d_t) + (1 + g(d_t))h_t & \\
b_{t+1} \geq 0 & \\
h_t^a \geq 0 & \\
h_t^d \geq 0 &
\end{aligned}$$

Households with children have a bequest motive in the last period (T), governed by the parameters κ and ϕ_{beq} .

$$\beta \kappa \ln(\phi_{\text{beq}} + b_{T+1} + h_T)$$

5.3 Total household consumption

Appendix C.1 shows the derivations of the first-order conditions for the problem above. By defining total household consumption of nondurables as $c_t = c_t^a + d_t c_t^d$ and of housing services as $m_t = m_t^a + d_t m_t^d$, the problems above can be written as in equations (9) and (10) (where $v(\cdot, \cdot, \cdot)$ is a function detailed at the end of Appendix C.1). This simplification will allow for an easier interpretation of the trade-offs faced by households in Section 5.4. Written this way, households choose total household consumption of nondurables and housing services, and its definition adjusts for the household composition, such as whether a child is living at home. This simplification allows us to see more clearly two effects of children. First, they increase the consumption of nondurables and housing services by $d_t c_t^d$ and $d_t m_t^d$, respectively. Second, they change the relative weights of these two goods. For example, in the case of nondurables, the weight for households without children is $\alpha_1/(\alpha_1 + \alpha_2)$, while for those with children it is

$(\alpha_1 + \psi_1 d_t)/(\alpha_1 + \psi_1 d_t + \alpha_2 + \psi_2 d_t)$. In Section 5.7, I discuss the magnitudes with the chosen parameter values.

For those who choose to rent, the problem in terms of total household consumption is

$$\begin{aligned}
V_t(b_t, h_{t-1}, y_t, d_t, f) &= \max_{c_t, b_{t+1}, l_t} \frac{v(c_t, m_t, d_t)^{1-\sigma}}{1-\sigma} + \beta \mathbb{E}_t[V_{t+1}(b_{t+1}, 0, y_{t+1}, d_{t+1}, f)] \quad (9) \\
&\text{s.t.} \\
&b_{t+1} + p_r l_t + c_t + \phi(h_{t-1}, 0) = R_b^f b_t + h_{t-1} + y_t \\
&m_t = l_t \\
&b_{t+1} \geq 0 \\
&l_t \geq 0.
\end{aligned}$$

For homeowners, the problem in terms of total household consumption is

$$\begin{aligned}
V_t(b_t, h_{t-1}, y_t, d_t, f) &= \max_{c_t, b_{t+1}, h_t} \frac{v(c_t, m_t, d_t)^{1-\sigma}}{1-\sigma} + \beta \mathbb{E}_t[V_{t+1}(b_{t+1}, h_t, y_{t+1}, d_{t+1}, f)] \quad (10) \\
&\text{s.t.} \\
&b_{t+1} + h_t + c_t + \phi(h_{t-1}, h_t) = R_b^f b_t + h_{t-1} + y_t \\
&m_t = \text{taste}_o(h_t, d_t) + (1 + g(d_t))h_t \\
&b_{t+1} \geq 0 \\
&h_t \geq 0.
\end{aligned}$$

5.4 Optimality conditions

For the purpose of the paper, the most interesting optimality conditions are those related to the choice between saving in the financial asset, b_{t+1} , and the durable good, h_t . From the solution in problem (10), we can derive two Euler conditions, presented in equations (11) and (12), where $v_{x,\tau}$ is the derivative of v with respect to x at time τ .

$$v_{c,t}^{-\sigma} = \beta R_b^f v_{c,t+1}^{-\sigma} \quad (11)$$

$$v_{c,t}^{-\sigma}(1 + \phi_2(h_{t-1}, h_t)) = v_{m,t}^{-\sigma}(1 + g(d_t)) + \beta(1 - \phi_1(h_t, h_{t+1}))v_{c,t+1}^{-\sigma} \quad (12)$$

The left-hand side of both equations denotes the cost of investing in these assets, which is related to nondurable consumption. The first difference here is the magnification of the cost of investing in h by the term $\phi_2(h_{t-1}, h_t)$, which denotes the adjustment cost of saving in housing. The right-hand side shows the benefits of saving in each asset. For the liquid asset, it depends on its return R_b^f , which varies by financial-literacy level. Abstracting from other effects, those with higher financial literacy prefer to invest in this asset as it has a higher return. However, the presence of children might change this. The benefits of saving in h are

the housing services in the same period and the availability of the asset to consume in the next period. The first term is affected by the composition of the household (d), which might make h more appealing.

These trade-offs can explain the patterns in the adjustment speed, especially in the middle and top of the wealth distribution, where these conditions probably hold with equality. When those in the middle of the wealth distribution have children, the household composition makes them prefer housing. In this context, those with higher financial literacy can overcome more easily the additional cost ($\phi_2(h_{t-1}, h_t)$), so they adjust housing before those with lower financial literacy. Those at the top, even with low financial literacy, probably have the resources to make the adjustment and own a big enough house.

For those at the bottom of the wealth distribution, some of the conditions (11) and (12) might hold with inequality. Those with higher financial literacy will have the condition in (11) hold with equality sooner, making them save in the financial asset. Later in their children's lives, they will also have the resources to save in housing since the household composition term might dominate.

5.5 Functional forms

The income process is composed of a deterministic part (y_t^f), a persistent z_t , and a transitory shock (ϵ_t).

$$\ln(y_t) = y_t^f + z_t + \epsilon_t \quad (13)$$

$$z_t = \rho_y z_{t-1} + \nu_t \quad (14)$$

The adjustment cost of housing is proportional to the value of the house bought in the case of a change of residence. As mentioned before, this is very standard in the literature (e.g., Bajari et al., 2013; Fernández-Villaverde and Krueger, 2011; Fella, 2014), and it captures real-estate fees paid by a real-estate buyer. Under this specification, the seller bears no transaction costs, meaning there are no adjustment costs from moving into renting.

$$\phi(h_{t-1}, h_t) = \begin{cases} 0 & h_t = h_{t-1} \\ \phi_1 h_t & h_t \neq h_{t-1} \end{cases} \quad (15)$$

The function $taste_o(h_t, d_t)$ is defined piecewise as an additional utility for owning differentiated by housing size. For the case of no children ($d_t = 0$),

$$taste_o(h_t, 0) = \begin{cases} taste_{0,1} & h_t > 0 \ \& \ h_t \leq \bar{h} \\ taste_{0,2} & h_t > \bar{h}. \end{cases} \quad (16)$$

For households with children ($d = 0.25$), it has a similar parametrization, with the same threshold (\bar{h}):

$$taste_o(h_t, 0.25) = \begin{cases} taste_{1,1} & h_t > 0 \ \& \ h_t \leq \bar{h} \\ taste_{1,2} & h_t > \bar{h}. \end{cases} \quad (17)$$

Function $g()$ is defined as $g(d) = \eta_{age}^d - 1$, for four age groups, the same used to document the stylized fact related to the adjustment speed in Section 4.3. The model only considers at most one child in each household ($d = 0.25$), so the parameter η_{age} governs the additional utility of housing for households with children at different ages.

$$g(d) = \begin{cases} 0 & d = 0 \\ \eta_{0-4}^d - 1 & d > 0, \text{Age} \leq 4 \\ \eta_{5-9}^d - 1 & d > 0, 5 \leq \text{Age} \leq 9 \\ \eta_{10-14}^d - 1 & d > 0, 10 \leq \text{Age} \leq 14 \\ \eta_{15-17}^d - 1 & d > 0, 15 \leq \text{Age} \leq 17 \end{cases} \quad (18)$$

5.6 Simulations

The estimation and welfare analysis comes from simulating 10,000 households for 23 periods, each corresponding to two years, from age 20 to age 66. They are born with a given financial-literacy level (high or low), a dummy for fertility, and the initial level of the permanent income shock, all three drawn from a copula reproducing correlation in the data. The moment of birth is either Period 4 or 7 and is known from moment 0. The initial assets are 60% of the initial income, corresponding to the median asset to income and the mean wealth to income in the data. The assets are housing with a probability of 13% or only deposits with the remaining probability, which corresponds to the probability in the data for those younger than 25. The probability of having an unplanned birth also depends on the initial level of the permanent income shock, estimated from the data. Those with the lowest level have a probability 0.4 of having an unplanned birth, those in the middle group have a probability of 0.2, and those with the highest level have a probability of 0.08.

Variables are in tens of thousands of 2018 US dollars. The variable h is partitioned into 21 points with a maximum of 60, while b is partitioned into 101 points with a maximum of 200. The permanent shocks, z_t , in three points and the transitory shock, ϵ , in two points (equations (13) and (14)). The model is solved using the endogenous-grid method developed by Fella (2014).

5.7 Parameter values

Most of the parameter values are externally calibrated, except for two sets that are very specific and relevant to the structural model: the return of the financial asset for both levels of financial literacy and parameters related to the preference of owning over renting. These parameters are key to reproducing the stylized facts presented in Section 4, as they govern the additional utility of housing and the additional return of the financial asset. Section 5.7.1 presents those parameters taken from the literature, while Section 5.7.2 explains the targeted moments and the values obtained for the other parameters.

5.7.1 External parameters

Table 6 shows the externally calibrated parameters. The values for the utility function are very standard. The elasticity of substitution between durable and nondurable goods comes from Fernández-Villaverde and Krueger (2011) and is equal to one ($\rho = 0$). This means that the CES becomes a Cobb-Douglas utility function. The authors find no clear consensus on this parameter, with estimations for ρ above and below 0. The weights in the CES (α_1, α_2) come from Kaplan and Violante (2014), who also include housing in their model and focus on the illiquidity of this asset. Similar values are found in the literature, which imply that the consumption of housing is 15% of total expenditure. The elasticity of intertemporal

Table 6: Externally calibrated parameter values

Parameter	Value	Source
ρ	0	Fernández-Villaverde and Krueger (2011)
α_1	0.85	Kaplan and Violante (2014)
α_2	0.15	Kaplan and Violante (2014)
σ	2.0	Fernández-Villaverde and Krueger (2011)
ψ_1	0.92	Attanasio et al. (2016)
ψ_2	0.92	Attanasio et al. (2016)
y_t^f		Hansen (1993)
ρ_y	0.935	Fernández-Villaverde and Krueger (2011)
σ_ν	0.247	Fernández-Villaverde and Krueger (2011)
σ_ϵ	0.130	Fernández-Villaverde and Krueger (2011)
κ	2.56	Bajari et al. (2013)
ϕ_{beq}	1.834	Cooper and Zhu (2016)
β	0.94	
g_f	0.0724	Bajari et al. (2013)
ϕ_1	0.06	Bajari et al. (2013)

substitution ($1/\sigma$) is set to 0.5, a standard value in the literature. For example, Kaplan and Violante (2014) uses $2/3$ for the same parameter.

The parameter ψ_1 is chosen so that the ratio of nondurable consumption of a household with and without a child matches the international estimations, as is standard in the literature, (particularly, Attanasio et al., 2016). In my case, the estimated parameter is almost double because households without children comprise a couple. Hence, a child represents 0.25 of my unit, not 0.5, when a child is compared to one adult. I assume the ratio is the same for durable goods (housing services), so $\psi_2 = \psi_1$. These parameter values imply that the expenditure on housing goes from 0.15 for households without children to 0.26 ($= (0.15 + 0.25 \times 0.92)/(1 + 0.5 \times 0.92)$) for those with children.

The parameters governing the income evolution come from Fernández-Villaverde and Krueger (2011). Following their choice, I take the deterministic component from Hansen (1993), a standard, inverted-U shape during the life cycle. The volatility and persistence of the income shocks are taken directly from Fernández-Villaverde and Krueger (2011). The bequest parameters are taken from Bajari et al. (2013) and Cooper and Zhu (2016), while the β implies an annual discount of 0.97, standard in the literature. The flow of services from the stock of housing in the utility function, g_f , comes directly from Bajari et al. (2013) as well as the adjustment cost, ϕ_1 , 6% (also used, e.g., by Fella, 2014).

5.7.2 Targeted moments

Table 7 presents the calibrated parameters for housing preferences and return of the financial asset and the corresponding moments used to calibrate them. The first three moments relate to the additional share of housing held by households with children in each tercile of the wealth distribution compared to households in the same tercile without children. The fourth moment is related to the aggregate return of the financial asset, which is set to 2% annually. The next four moments capture the difference in housing share between parents with high and low financial literacy for different ages of their children, as shown in Figure 11b. Finally, three moments capture the share of owners in each tercile of the income distribution.

The parameters to calibrate with these moments are the returns of financial literacy ($R_{\text{lowfinlit}}$ and $R_{\text{highfinlit}}$), the parameters of the housing preferences for households without children ($taste_{0,1}$, $taste_{0,2}$, and \bar{h}), and the parameters related to the housing preferences of households with children ($taste_{0,1}$, $taste_{0,2}$, η_{0-4} , η_{5-9} , η_{10-14} , and η_{15-17}). The number of parameters and moments are equal, so the model is just identified.

Although all the parameters are jointly estimated, some moments more related to some parameters. For example, the moments of share of owners without children identify the parameters $taste_{0,1}$, $taste_{0,2}$, and \bar{h} . The moments from Figure 11b identify the parameters η , which vary by group age. Finally, the aggregate return on the financial asset and parents'

Table 7: Targeted moments and parameter values

Parameter	Value	Moment	Model	Data
η_{0-4}	0.80	Child at home tercile 1	0.123	0.125
η_{5-9}	0.52	Child at home tercile 2	0.024	0.084
η_{10-14}	0.41	Child at home tercile 3	0.088	0.062
η_{15-17}	0.37	R	1.037	1.040
$R_{\text{low fin. lit.}}$	1.027	Figure 11b 0–4	0.046	0.047
$R_{\text{high fin. lit.}}$	1.049	Figure 11b 5–9	0.100	0.097
$taste_{0,1}$	3.88	Figure 11b 10–14	0.078	0.056
$taste_{0,2}$	5.18	Figure 11b 15–17	−0.004	−0.012
$taste_{1,1}$	7.49	Share owners ($d = 0$) age 40 tercile 1	0.096	0.185
$taste_{1,2}$	7.67	Share owners ($d = 0$) age 40 tercile 2	0.248	0.435
\bar{h}	6.69	Share owners ($d = 0$) age 40 tercile 3	0.679	0.681

The first three moments are the excess of share in housing for households with and without children. Moments related to Figure 11b are the point estimates for different ages. Terciles of owners refers to income terciles.

additional share of housing for different terciles identify the parameters on the additional utility of parents and the two returns in the financial asset.

The model can match the empirical moments, especially those related to the adjustment speed in the middle of the wealth distribution, where households with higher financial literacy adjust their portfolios towards housing earlier in their children’s lives. It also matches the aggregate return and the share of households without children who own housing, which is increasing in income.

Some implications can be compared despite the parameters not being comparable to other models. Regarding the returns, the gap between high and low financial literacy is around 1% annually. This seems reasonable when compared to empirical estimations, such as Bianchi (2018), who estimates 0.4% after adjusting for risk. The difference in the composition of safe financial assets can explain the remaining difference. As explained in Section 2.3, the empirical counterpart of the financial asset is composed not only of checking and savings accounts, but also of fixed-term deposits and other financial assets. Thus, the difference in return for the two levels of financial literacy can also be explained by heterogeneity in the composition of their financial asset, where those with higher financial literacy own assets with higher returns and are able to obtain a higher return in the same asset class.

The parameters related to housing preferences cannot be compared, but the average preference for owning can be compared to other papers in the literature. The implied average ratio in housing services (m) from owning and renting in the model is 1.3. This is a very




Figure 12: Share of housing. Double difference by fertility and financial literacy: comparison between financially literate and illiterate parents in the same wealth tercile. The plots replicate Figures 11c and 11a from the data

similar number to the high value of the shock estimated by Bajari et al. (2013). Moreover, Bacher (2021) finds that owning the smallest house yields a utility four times larger than renting it. In my case, those without children have twice as much, and those with children three times as much.

5.7.3 Nontargeted patterns

The model can correctly replicate two sets of empirical patterns, which are central to the paper. First, the adjustment speed for the bottom and top of the wealth distribution. The calibration process used that in the middle of the distribution (Figure 11b). The patterns in Figures 11a and 11c, which were not used in the calibration process, are replicated by the model.

Figure 12a shows how, for richer households, the model predicts almost no difference in the adjustment speed towards housing when comparing financially literate and illiterate parents. It also replicates the nonsignificant downward trend in the children’s lives. More critically, the model is able to reproduce the delay at the bottom of the wealth distribution. Figure 12b shows how, in the model, financially literate parents at the bottom of the wealth distribution delay the adjustment towards housing compared to the financially illiterate. Before children turn 10, the former have a lower share of housing in their portfolio, but this pattern reverses after their children turn 10, when they hold a higher share of housing.

The second pattern the model can reproduce is the decrease in the share of housing after a shock (Figure 2). The model is able to reproduce this decrease in all wealth terciles. Figure 13 shows the change in the share of housing when comparing households who faced a fertility shock and those who did not have children. The model’s predictions are always




Figure 13: Comparison between model and data. Share of housing around a shock. Control group: households with no birth. Before birth share normalized to 0.

within the confidence intervals, and the direction is correct in the three terciles. The model also replicates the inverted-U shape, in which those in the middle of the wealth distribution decrease their share of housing the least.

One disadvantage of the data is that fertility shocks and financial-literacy levels come from different datasets, making it impossible to empirically find differences in the adjustment for households with different financial-literacy levels. This is one of the purposes of building the model. The next section clears this difficulty and discusses the mechanism behind it.

5.8 Shocks and portfolio adjustments

Households with children prefer owning a larger share of housing in their portfolios, as documented in Section 4.2. The model replicates these patterns with the functions $taste(h, d)$ and $g(d)$. However, this adjustment has a cost, represented in the model by the function $\phi(h_{t-1}, h_t)$, particularly the parameter ϕ_1 . Therefore, even those households with an unplanned birth want to own a house, yet the difference with those who had planned births is the lack of planning in this decision. A shock forces households to save in the financial asset to have more liquidity to face the larger expenditures and to be able to overcome the adjustment cost of owning in the future. One can imagine how those with a higher return on this asset can meet the increase in expenditures more easily and pay for the adjustment cost for owning, probably for a larger house, as they can save more. This mechanism should lead households with higher financial literacy to smooth the shock more easily and behave more as those who did not face a shock.

In the model, this implies that households with higher financial literacy should reduce the share of housing in their portfolios to a *lesser extent* compared to those with lower financial literacy. Figure 14 illustrates this phenomenon across all wealth levels; blue dots represent

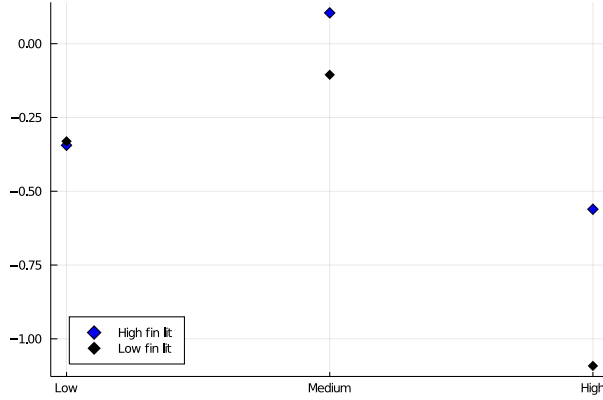


Figure 14: Share of housing around a shock. Control group: households with no birth. Before-birth share normalized to 0.

high-financial-literacy households and gray dots represent low-financial-literacy households. For poorer households, the difference is almost zero. Most of these households are constrained and cannot use the financial asset to save. Therefore, even having access to higher returns makes little difference for parents facing a fertility shock. This underscores a crucial policy implication: Financial literacy serves as an enabler for savings and portfolio management only when households have savings to leverage this knowledge.

The middle of the wealth-distribution results emphasize the disparity in portfolio allocation based on financial-literacy levels. Those in the second tercile with high financial literacy can even increase the share of housing after a shock, akin to households with planned births. In contrast, households in the second tercile with low financial literacy are forced to decrease the share of housing. Financially literate households can accommodate additional expenditures due to the interest they receive from the financial asset, making it easier to overcome adjustment costs. In contrast, financially illiterate parents must increase their savings in financial assets to meet the expenditure increase. The top of the wealth distribution also shows a notable difference in portfolio allocation, with financially literate individuals experiencing a less significant decrease in the share of housing.

6 Welfare: Financial literacy and fertility

One of the model's advantages is its ability to compare the welfare of groups with different characteristics. All the results are derived from the life-cycle model, which is a partial-equilibrium model. Some magnitudes might be smaller in a general-equilibrium model, but the results are large enough to expect they will remain directionally true in a general-equilibrium setting. Since the model includes children in the utility function, all comparisons are between groups with the same fertility; households with children are compared among

themselves and those without similarly.

The welfare measure used is in terms of units of the nondurable good (c). When comparing groups i and j , high and low financial literacy, for example, I take the average utility among both and find the value m that matches both while keeping other variables constant, as shown in equation (19).

$$\mathbb{E} \sum_{t=0}^T \beta^t u(c_t^1, h_t^1, d_t^1) = \mathbb{E} \sum_{t=0}^T \beta^t u(m \times c_t^2, h_t^2, d_t^2) \quad (19)$$

This section presents two main comparisons. First, Section 6.1 presents the welfare comparison between groups with low and high financial literacy for different cases of fertility. That is, how much does financial literacy increase the welfare for those without children, those with planned children, and those with unplanned children. Section 6.2 shows the welfare comparison between planned and unplanned children by level of financial literacy and position in the wealth distribution. This allows us to analyze whether financial literacy can help dampen the negative effects of shocks.

6.1 Welfare: Fertility and financial literacy

This section highlights the importance of the interaction between financial literacy and fertility, a relationship not analyzed in the literature. This is another example of the importance of financial literacy, not only for managing financial assets, but for improving the welfare of groups with more expenditures and a greater need for resources.

Table 8 shows that in households with children, financial literacy improves welfare even more. In households without children, taken as the benchmark, financial literacy increases welfare by 18%. This is large, but it is explained by a larger return on the financial asset and the fact that these are partial-equilibrium results. For households with planned children, this number goes to 28%, 1.5 times the benchmark. This result is explained by parents' preference to own housing. This pushes them to save more when their children are young by buying housing. Thus, they have more savings to consume later in their lives. Households

Table 8: Welfare gain by financial literacy

Group	Welfare
No child. High vs low fin. lit.	0.18
Planned child. High vs Low fin. lit.	0.28
Unplanned. High vs low fin. lit.	0.22

Comparison in nondurable consumption units, equation (19).

with unplanned children also have a larger welfare gain by having high financial literacy, 22%. The mechanism affecting the households with planned children is also present in this group. However, since they have to adjust to a shock, they lose some of that gain. As shown in Section 5.8, they decrease the share of housing in the first years after having a child, dampening the accumulation channel of wealth in housing.

6.2 Welfare: Financial literacy and shocks

This section evaluates how financial literacy can dampen the welfare costs of shocks. I focus on the first two terciles of the wealth distribution, as these are the households that might be constrained and for whom financial literacy can be more valuable.³ Table 9 presents the welfare losses of having an unplanned child compared to having a planned one for a given wealth and financial-literacy level. Although a more precise comparison would be between those with unplanned children and those without, as this was the expected path for them before the shock, including children in the utility function makes it unfeasible.

From Table 9 at the bottom of the wealth distribution, those with low financial literacy have a welfare loss of 3% when they have an unplanned child compared to planing couples. In contrast, those with higher financial literacy only have a welfare loss of 2.3%. This means that financial literacy protects in some sense from the fertility shock, as the welfare loss is smaller. In particular, it is 23% smaller in this case($2.3\%/3.0\% - 1$). For the middle of the wealth distribution, this value is higher. As shown in Section 5.8, this group’s portfolio-adjustment gap is larger. This is reflected in a higher compensation of financial literacy to welfare losses. In fact, in this part of the wealth distribution, this protection is almost 100%, as the welfare loss of unplanned children is almost 0.0% for those with high financial literacy.

These results have very important policy implications. The possible increase of unplanned children by the overturn of Roe v Wade will pose challenges for many households, particularly those at the bottom or in the middle of the wealth distribution, who are more exposed, as shown in Figure 1b. Financial literacy could help these households, giving them appropriate

³Results for the top tercile are presented in Appendix C.2.

Table 9: Welfare loss by a shock

Group	Low Wealth	Med. Wealth
Low fin. lit.	-0.030	-0.070
High fin. lit.	-0.023	-0.004

Planned vs unplanned births by wealth and financial literacy level. Comparison in nondurable consumption units, equation (19).

tools to manage their finances in these cases.

7 Conclusions

Financial literacy plays an important role by helping households face shocks. This paper's structural model showed how more financially literate households with fertility shocks can have portfolio adjustments closer to those with planned children. This results in financial literacy dampening welfare losses for more vulnerable households. Health shocks can be similar in nature as households also face increased expenditures. The ability to adjust the liquidity of their portfolio accordingly also plays a crucial role in helping households face the expenditure increases.

My empirical analysis shows that households tend to accumulate more housing wealth around a child's birth, and that financial literacy significantly impacts the portfolio-adjustment speed. For fertility shocks, the empirical results show a decrease in the share of housing, consistent with households increasing their portfolio liquidity to face the expenditure increases.

One of this paper's principal contributions is the structural model developed to analyze the mechanism and counterfactuals. The model is parsimonious, yet it replicates the empirical patterns and facilitates the mechanism analysis. It highlights that financial literacy helps households overcome transaction costs associated with portfolio adjustments while maintaining an adequate liquidity buffer. Welfare analysis highlights the importance of financial literacy in portfolio adjustments and its importance in welfare gains. Results point to financial competence compensating at least 20% of the welfare losses.

Policymakers should take into account these benefits when considering financial-education programs. Using a structural model, this is one of the first papers to show how financial literacy plays a greater role in households with higher expenditures who might face liquidity constraints. Financial knowledge is important not only for those who own financial assets, but also for those with limited resources who need to increase their savings returns to face shocks more easily. The effect is greater for those with more savings, who have the opportunity to employ their knowledge. This proves to be an important interaction when considering financial-literacy programs: Resources enhance the efficacy of financial literacy.

Future research should analyze how financial literacy can affect consumption decisions, helping households understand their financial implications. This new channel of financial literacy would be relevant not only for fertility choices, but for other persistent expenditure decisions, such as rent or mortgages. Future research should also address the general-equilibrium implications of financial-literacy programs. Increasing financial literacy will likely generate more returns for households who are now financially illiterate, but it might also decrease the returns for those who are now financially literate.

References

- Angrisani, Marco et al. (2020). *The Stability and Predictive Power of Financial Literacy: Evidence from Longitudinal Data*. NBER Working Papers 28125. National Bureau of Economic Research, Inc.
- Attanasio, O. et al. (2016). “Global Demographic Trends”. In: *Handbook of the Economics of Population Aging*. Ed. by John Piggott and Alan Woodland. Vol. 1. Handbook of the Economics of Population Aging. Elsevier. Chap. 0, pp. 179–235.
- Babiarz, Patryk and Cliff Robb (2014). “Financial Literacy and Emergency Saving”. In: *Journal of Family and Economic Issues* 35(1), pp. 40–50.
- Bacher, Annika (2021). *Housing and savings behavior across family types*. URL: https://annikabacher.github.io/Bacher_JMP.pdf.
- Bajari, Patrick et al. (2013). “A Dynamic Model of Housing Demand: Estimation and Policy Implications”. In: *International Economic Review* 54(2), pp. 409–442.
- Baker, Scott R. et al. (2023). “Income, Liquidity, and the Consumption Response to the 2020 Economic Stimulus Payments”. In: *Review of Finance*, art. rfad010.
- Bar, Michael et al. (2018). “Why Did Rich Families Increase Their Fertility? Inequality and Marketization of Child Care”. In: *Journal of Economic Growth* 23(4), pp. 427–463.
- Bianchi, Milo (2018). “Financial Literacy and Portfolio Dynamics”. In: *Journal of Finance* 73(2), pp. 831–859.
- Bogan, Vicki L. (2013). “Household Investment Decisions and Offspring Gender: Parental Accounting”. In: *Applied Economics* 45(31), pp. 4429–4442.
- Coile, Courtney and Kevin Milligan (2009). “How Household Portfolios Evolve After Retirement: The Effect of Aging and Health Shocks”. In: *Review of Income and Wealth* 55(2), pp. 226–248.
- Cooper, Russell and Guozhong Zhu (2016). “Household Finance over the Life-Cycle: What Does Education Contribute?” In: *Review of Economic Dynamics* 20, pp. 63–89.
- De Ferra, Sergio, Kurt Mitman, and Federica Romei (2020). “Household Heterogeneity and the Transmission of Foreign Shocks”. In: *Journal of International Economics* 124(C), art. 103303.
- Deuffhard, Florian, Dimitris Georgarakos, and Roman Inderst (2019). “Financial Literacy and Savings Account Returns”. In: *Journal of the European Economic Association* 17(1), pp. 131–164.
- Fagereng, Andreas et al. (2020). “Heterogeneity and Persistence in Returns to Wealth”. In: *Econometrica* 88(1), pp. 115–170.
- Fella, Giulio (2014). “A Generalized Endogenous Grid Method for Non-smooth and Non-concave Problems”. In: *Review of Economic Dynamics* 17(2), pp. 329–344.

- Fernández-Villaverde, Jesús and Dirk Krueger (2011). “Consumption and Saving Over the Life Cycle: How Important Are Consumer Durables?” In: *Macroeconomic Dynamics* 15(5), pp. 725–770.
- Gallipoli, Giovanni and Sebastian Gomez-Cardona (Feb. 2023). *The Production of Financial Literacy*. Working Papers 2023-007. Human Capital and Economic Opportunity Working Group.
- Gaudecker, Hans-Martin von (2015). “How Does Household Portfolio Diversification Vary with Financial Literacy and Financial Advice?” In: *The Journal of Finance* 70(2), pp. 489–507.
- Görlitz, Katja and Marcus Tamm (2015). *Parenthood and Risk Preferences*. Ruhr Economic Papers 552. RWI - Leibniz-Institut für Wirtschaftsforschung. Ruhr-University Bochum. TU Dortmund University. University of Duisburg-Essen.
- Hansen, G. D. (1993). “The Cyclical and Secular Behaviour of the Labour Input: Comparing Efficiency Units and Hours Worked”. In: *Journal of Applied Econometrics* 8(1), pp. 71–80.
- Hastings, Justine and Olivia S. Mitchell (2020). “How Financial Literacy and Impatience Shape Retirement Wealth and Investment Behaviors”. In: *Journal of Pension Economics and Finance* 19(1), pp. 1–20.
- Kaplan, Greg and Giovanni L. Violante (2014). “A Model of the Consumption Response to Fiscal Stimulus Payments”. In: *Econometrica* 82(4), pp. 1199–1239.
- Kaplan, Greg and Giovanni L. Violante (2018). “Microeconomic Heterogeneity and Macroeconomic Shocks”. In: *Journal of Economic Perspectives* 32(3), pp. 167–194.
- Klapper, Leora, Annamaria Lusardi, and Georgios A. Panos (2013). “Financial Literacy and Its Consequences: Evidence from Russia During the Financial Crisis”. In: *Journal of Banking & Finance* 37(10), pp. 3904–3923.
- Liu, Kai (2016). “Insuring Against Health Shocks: Health Insurance and Household Choices”. In: *Journal of Health Economics* 46, pp. 16–32.
- Love, David A. (2009). “The Effects of Marital Status and Children on Savings and Portfolio Choice”. In: *The Review of Financial Studies* 23(1), pp. 385–432.
- Luetticke, Ralph (2021). “Transmission of Monetary Policy with Heterogeneity in Household Portfolios”. In: *American Economic Journal: Macroeconomics* 13(2), pp. 1–25.
- Lusardi, Annamaria and Carlo de Bassa Scheresberg (2013). *Financial Literacy and High-Cost Borrowing in the United States*. NBER Working Papers 18969. National Bureau of Economic Research, Inc.
- Lusardi, Annamaria, Andrea Hasler, and Paul J. Yakoboski (2021). “Building up Financial Literacy and Financial Resilience”. In: *Mind & Society: Cognitive Studies in Economics and Social Sciences* 20(2), pp. 181–187.

- Lusardi, Annamaria, Pierre-Carl Michaud, and Olivia S. Mitchell (2017). “Optimal Financial Knowledge and Wealth Inequality”. In: *Journal of Political Economy* 125(2), pp. 431–477.
- Lusardi, Annamaria and Olivia S. Mitchell (2007). “Financial Literacy and Retirement Preparedness: Evidence and Implications for Financial Education”. In: *Business Economics* 42(2), pp. 35–44.
- Miller, Amalia (2011). “The Effects of Motherhood Timing on Career Path”. In: *Journal of Population Economics* 24(3), pp. 1071–1100.
- Su, Jessica Houston and Fenaba R. Addo (2018). “Born Without a Silver Spoon: Race, Wealth, and Unintended Childbearing”. In: *Journal of Family and Economic Issues* 39(4), pp. 600–615.
- Urban, Carly et al. (2020). “The Effects of High School Personal Financial Education Policies on Financial Behavior”. In: *Economics of Education Review* 78, art. 101786.
- Van Rooij, Maarten C. J., Annamaria Lusardi, and Rob J. M. Alessie (2012). “Financial Literacy, Retirement Planning and Household Wealth”. In: *The Economic Journal* 122(560), pp. 449–478.

Appendix A Data

A.1 The Big 3—Financial-literacy questions

The Big 3 questions, originally developed by Lusardi and Mitchell (2007), provide a measure of financial literacy with three simple questions about inflation, diversification, and compound interest rates.

1. Imagine that the interest rate on your savings account was 1% per year and inflation was 2% per year. After 1 year, how much would you be able to buy with the money in this account?
 - (a) More than today
 - (b) Exactly the same
 - (c) **Less than today**
 - (d) Don't know
 - (e) Prefer not to say

2. Buying a single company's stock usually provides a safer return than a stock mutual fund.
 - (a) True
 - (b) **False**
 - (c) Don't know
 - (d) Prefer not to say

3. Suppose you had \$100 in a savings account, and the interest rate was 2% per year. After 5 years, how much do you think you would have in the account if you left the money to grow?
 - (a) **More than \$102**
 - (b) Exactly \$102
 - (c) Less than \$102
 - (d) Don't know
 - (e) Prefer not to say

A.2 Questions on unplanned births

The question used in the benchmark definition of “unplanned” birth from the PSID is CAH100, a question asked to mothers for births after 2013. For those who answered “no,” the birth is classified as “unplanned.”

1. Just before you became pregnant (with CHILD NAME, or NAMES if multiple births), did you yourself want to have a baby?
 - Yes
 - No
 - DK; NA; RF

I consider this a definition restrictive enough to identify fertility shocks while keeping a large sample size. The PSID includes further questions related to the timing of the child—whether the mother did not want to have a child at that particular moment—and the same set of questions for the father.

A.3 Other definitions of “unplanned” children

The alternate definition uses questions on contraceptive methods. Those mothers who answered that they were using contraceptive methods before the pregnancy and did not stop them before getting pregnant are classified as fertility shocks. The specific questions are CAH96 and CAH97 from the Childbirth and Adoption History supplement.

1. Just before you became pregnant with (CHILD NAME, or NAMES if multiple births), did you use any methods to keep from getting pregnant?
 - Yes
 - No
 - Never used any methods to prevent pregnancy
 - DK; NA; RF
2. Had you stopped all methods before you became pregnant?
 - Yes
 - No
 - Never used any methods to prevent pregnancy
 - DK; NA; RF

Figure 15 shows the share of births classified as “unplanned” using these questions. It follows the same decreasing trend as Figure 1b, which uses the benchmark definition. The proportion is much lower than the benchmark. As noted in the main text, only 70 births are classified as “unplanned” with this definition. Interestingly, the major difference between the two definitions comes from women who did not use contraceptives and did not want to have children. This could be an interesting avenue of future research for household composition. The benchmark definition is a good proxy of “unplanned” births, as it measures how expected it was. This answer is also informative of how prepared the household was in another aspect, such as finances.

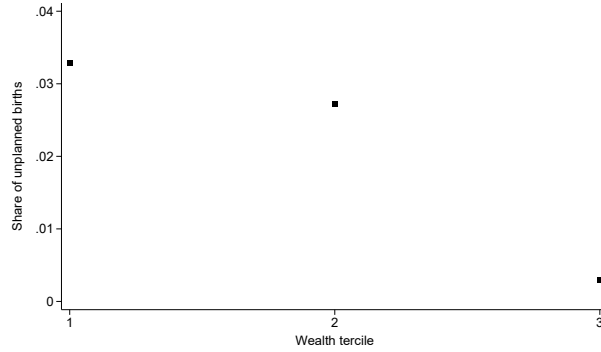


Figure 15: Share of “unplanned” births by wealth tercile according to definition using contraception questions. Source: PSID

A.4 Data selection

In the PSID, I restrict the sample to households with a head between 20 and 65 years old. Households with any major composition change in their members are dropped. I keep households with no change; the change was in other than the reference, the spouse, the spouse died, but the rest of the members are the same; or when the spouse became the reference person.

From the SCF, I only use the waves in 2016 and 2019, as these are the ones that include the financial-literacy questions. I restrict the sample to households with heads between 20 and 65. The SCF includes five implicates of the data. I only use Implicate 3. The results are robust to the choice of the implicate.

Appendix B Robustness checks—Empirical results

B.1 Effect of unplanned births on other assets

From the assets included in the household’s portfolios—safe financial assets, main residence, stocks, retirement accounts, other real estate, and private businesses—the largest negative effect is on housing, while the only positive one is on the safe financial asset. Moreover, the two assets in the paper, safe financial assets and main residence have the largest difference between before and after the shock. For completeness, Figure 16 presents the results for the other assets. Stocks and retirement accounts also show negative effects after a shock. I exclude these assets because they are not as relevant for poor households. In future research, the implications for retirement should be analyzed, particularly the role of financial literacy, which has been proven to be a key factor in retirement preparedness (e.g., Lusardi and Mitchell, 2007; Van Rooij, Lusardi, and Alessie, 2012; Hastings and Mitchell, 2020).

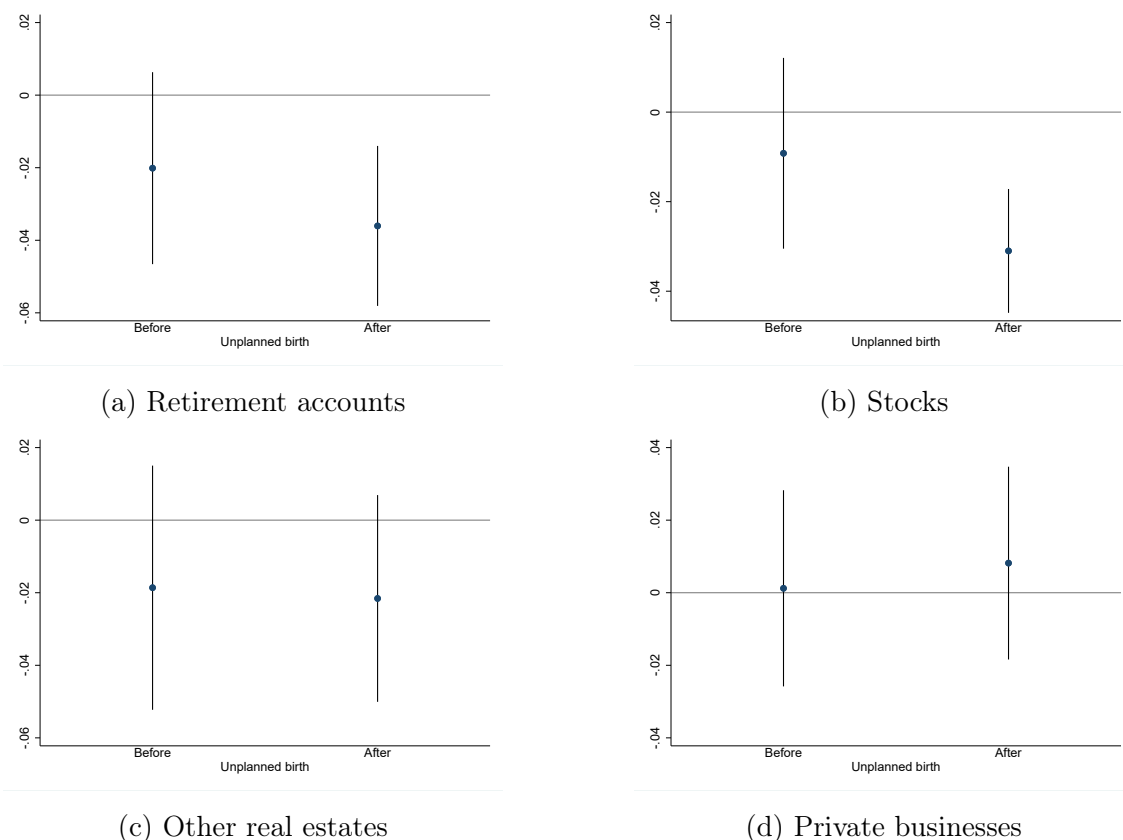


Figure 16: Share before and after a fertility shock. Control group: households without children. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.

B.2 Fertility shocks—Results with alternative definition

Figure 17 compares the share of housing and the safe financial asset for the two definitions of fertility shocks. The benchmark uses the question of whether the child was wanted, while the contraception uses the question of contraceptive methods. In both cases, the share of housing decreases after the shock, while the share of the safe financial asset increases. The choice of the benchmark is mostly about sample size, which allows for analysis of the effect of wealth tercile.

B.3 Fertility shocks—Portfolio path

The analysis around fertility shocks can be decomposed observation by observation instead of grouping three observations before and after the shock. Figure 18 shows how, at moment t , which is the moment of the shock, the share of the safe financial asset jumps and that in




Figure 17: Share before and after a fertility shock. Comparison between benchmark definition of shock and contraception use as explained in Section A.3. Control group: households without children. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.




Figure 18: Share around a fertility shock. Unplanned birth at t . Control group: households without children. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.

housing decreases. Since the portfolio adjustment could take time, particularly in the presence of illiquid assets, the *before* and *after* analysis could be more informative, especially with survey data, which can be noisy.

B.4 Fertility shocks—Control group planned births

The effect of an unplanned child can also be compared to those with children. The direction of the effect is the same, as seen in Figure 19. While the share of housing decreases, that of the safe financial asset increases. As explained above, while those with planned children increase




Figure 19: Share before and after a fertility shock. Control group: households with planned births. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.

the share of housing after childbirth, those with unplanned children decrease it. Thus, the effect when the control group is those with children is larger. While it is not statistically significant, the decrease in housing when the control group is those without children is 4.6% against 5.1% when compared to those with planned children.

B.5 Family size—Number of children

The number of children plays a role in the parents' portfolio adjustments. In the empirical section and the model, I abstract from the number of children to make results more tractable. When including this additional variable in the analysis, a U-shape pattern is present in the results. Households with two children present the largest adjustments in the financial asset and housing (Figure 20). In my sample, the average number of children by household is close to 2, so the results of dummies for children at home probably pick up a number close to the effect for households with two children.

As mentioned in the paper's conclusions, future research should include the number of children as a margin in which households can decide and adjust their portfolios. Financial literacy could play a role from a different perspective: committed expenditures. The patterns in Figure 20 show some increase in the housing share for households with children. Those with two children have a significantly higher share than those with one. This comes with a lower share of the safe financial asset. However, those with three children or more revert to this pattern, with shares in housing closer to those with one child. More children could constrain households into adjusting their portfolios in the desired direction. A larger panel, with more years of information about unplanned children, could shed some light on whether




Figure 20: Share in portfolio. Control group: households without children. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tercile, and wealth tercile.

families with more than two children are more likely to have unplanned children, which could explain the patterns from Figure 20.

B.6 Planning horizon—Married couples

Section 4.2.2 presents the housing path in parents' portfolios around children's births. Figure 6a shows how parents increase this share before childbirth. The graph has some upward trend, which could be a pretrend due to other circumstances, for example, and increase in housing due to marriage. As pointed out by Bacher (2021), marriage is an important determinant of housing demand. Figure 6a controls for marital status, but further restricting the sample could shed some light if children do increase housing in portfolios.

Figure 21 replicates the regression from Figure 6a but restricts the sample to married households and adds variables to control the number of years the head of the household has been married. The results are not contaminated by any effect that marriage can have on housing, not only by restricting the sample but also by controlling for the time of marriage, which takes into account the path of housing share after marrying, as documented by Bacher (2021).

The pattern in Figure 21 shows that even among those married, children are related to an increase in the share of housing. Although the pretrend is not eliminated, it is flatter. This sample also includes a planning horizon of one period, which is two years, showing that parents plan the birth of children and adjust their portfolio. The overall adjustment is larger. When comparing $t - 2$, four years before the birth, to t , the moment of childbirth, the difference in Figure 6a is 0.05 and in Figure 21 0.08. Therefore, controlling for marital

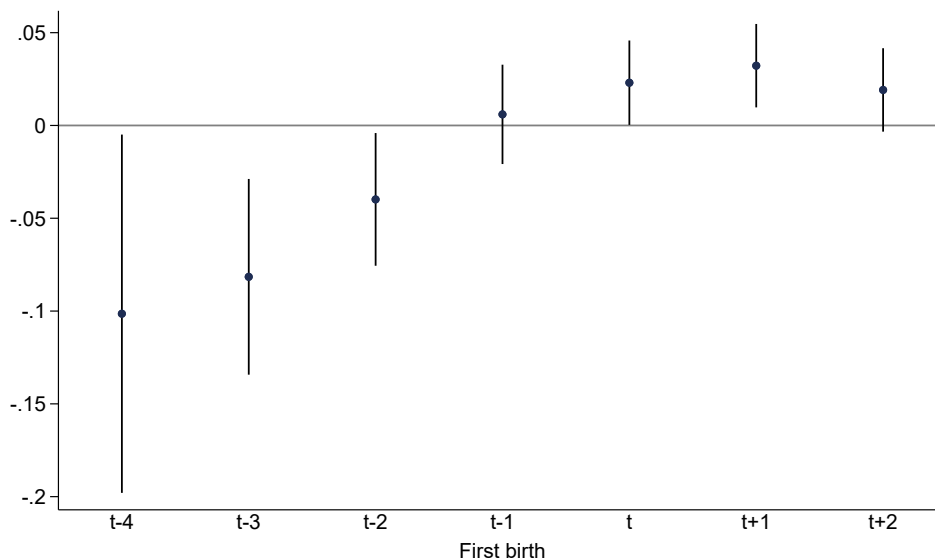


Figure 21: Share in portfolio. Control group: married households without children. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, and gender. Controls on household: state, number of adults, year fixed effects, income tercile, wealth tercile, and years married (when available).

status, there is still an increase in housing by parents around childbirth.

B.7 Adjustment speed—Extensive margin

The difference in the adjustment speed between parents with high and low financial literacy (Figure 11) is not only due to some parents buying bigger houses, but also due to some parents who previously rented buying houses. As shown in several adjustments, the extensive margin plays an important role. Figure 22 shows the speed of adjustment for the extensive margin, which follows the same pattern as the overall adjustment.

Interesting patterns emerge from the magnitudes of the extensive margin by taking the difference between the overall adjustment (Figure 11) and that in the extensive margin (Figure 22). Low in the wealth distribution, the absolute magnitude of the extensive margin's delay is larger, so the intensive margin adjustment is positive, implying that high-financial-literacy households, who own housing when they have a child, increase their share of housing, possibly by buying a larger house. The opposite pattern repeats when children approach adulthood, where the intensive margin seems negative. Putting these two together, the intensive margin adjustment behaves as in the top of the wealth distribution (Figure 11c).

In the middle of the wealth distribution (Figures 11b and 22b), a similar pattern occurs. Those with higher financial literacy adjust positively in the intensive margin and negatively

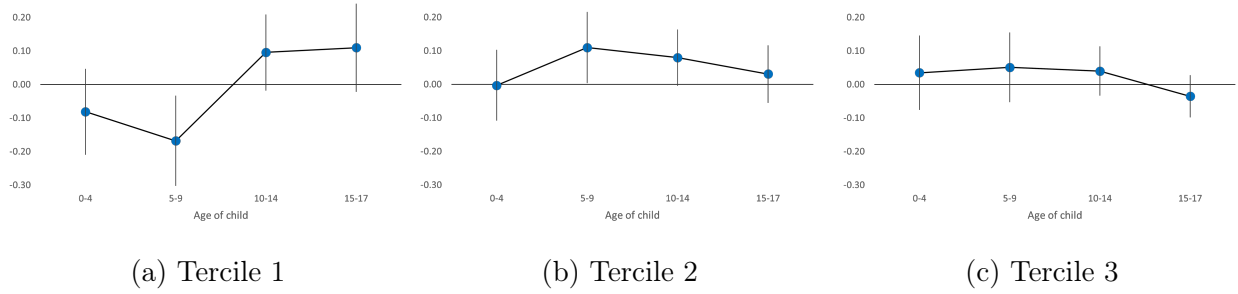


Figure 22: Effect of probability of owning housing. Double difference by fertility and financial literacy by wealth tertile. Calculated using *margins* in Stata. Controls on head of household: age, age squared, years of education, gender, and marital status. Controls on household: state, number of adults, year fixed effects, income tertile, and wealth tertile.

when children reach adulthood. For the top tertile, both effects are null as wealthy households can compensate for financial illiteracy.

Appendix C Model

C.1 Model: First-order conditions

To gain some intuition about the optimal choices, assume that both l and h are continuous variables, so we can interpret the derivatives. To simplify the notation, let $u(c, m, d)$ be the instantaneous utility function. That is,

$$u(c_t, m_t, d_t) = \frac{(\alpha_1(c_t^a)^\rho + \psi_1 d_t (c_t^d)^\rho + \alpha_2 (g_f m_t^a)^\rho + \psi_2 d_t (g_f m_t^d)^\rho)^{(1-\sigma)/\rho}}{1 - \sigma}. \quad (20)$$

Let λ_t^z denote the Lagrange multiplier of the budget constraint for either renting ($z = l$) or owning renting ($z = h$). In the case of renting, the first-order conditions, omitting any binding by the borrowing constraints, are

$$[c_t^a] : u_{c_t^a} - \lambda_t^l = 0 \quad (21)$$

$$[c_t^d] : u_{c_t^d} - d_t \lambda_t^l = 0 \quad (22)$$

$$[l_t^a] : u_{m_t^a} - p_r \lambda_t^l = 0 \quad (23)$$

$$[l_t^d] : u_{m_t^d} - p_r d_t \lambda_t^l = 0 \quad (24)$$

$$[b_{t+1}] : -\lambda_t^l + R_b^f \lambda_{t+1}^l = 0. \quad (25)$$

For the case of owning, the first-order conditions are

$$[c_t^a] : u_{c_t^a} - \lambda_t^h = 0 \quad (26)$$

$$[c_t^d] : u_{c_t^d} - d_t \lambda_t^h = 0 \quad (27)$$

$$[h_t^a] : u_{m_t^a}(1 + g(d_t)) - \lambda_t^h(1 + \phi_2(h_{t-1}, h_t)) - \lambda_{t+1}^h(\phi_1(h_t, h_{t+1}) - 1) = 0 \quad (28)$$

$$[h_t^d] : u_{m_t^d}(1 + g(d_t)) - d_t \lambda_t^h(1 + \phi_2(h_{t-1}, h_t)) - d_t \lambda_{t+1}^h(\phi_1(h_t, h_{t+1}) - 1) = 0 \quad (29)$$

$$[b_{t+1}] : -\lambda_t^h + R_b^f \lambda_{t+1}^h = 0. \quad (30)$$

From equations (21) and (22) one can get the relationship between c^a and c^d (which is the same as the one derived from equations (26) and (27)):

$$\alpha_1 (c_t^a)^{\rho-1} d_t = (\psi_1 d_t) (c_t^d)^{\rho-1} \Rightarrow c_t^d = (\alpha_1 / \psi_1)^{1/(\rho-1)} c_t^a. \quad (31)$$

Defining total household consumption as $c_t = c_t^a + d_t c_t^d$ one gets to the expression $c_t = c_t^a (1 + d_t (\alpha_1 / \psi_1)^{1/(\rho-1)})$ so that the expression inside the utility function becomes

$$\alpha_1 (c_t^a)^\rho + \psi_1 d_t (c_t^d)^\rho = \alpha_1 \left(1 + d_t \left(\frac{\alpha_1}{\psi_1} \right)^{\frac{1}{\rho-1}} \right)^{-\rho} c_t^\rho + \psi_1 d_t \left(1 + d_t \left(\frac{\alpha_1}{\psi_1} \right)^{\frac{1}{\rho-1}} \right)^{-\rho} \left(\frac{\alpha_1}{\psi_1} \right)^{-\frac{\rho}{\rho-1}} c_t^\rho \quad (32)$$

For the particular case considered in the paper, in which $\rho = 0$, the multiplicative term becomes $\alpha_1 + \psi_1 d_t$. Similarly, defining total household housing services as $m_t = m_t^a + d_t m_t^d$, one can prove that the housing-services terms of the utility function becomes a multiplicative term which reduces to $\alpha_2 + \psi_2 d_t$ for the case of $\rho = 0$. In this case, the CES part of the utility function reduces to a Cobb–Douglas with weight $(\alpha_1 + \psi_1 d_t) / (\alpha_1 + \psi_1 d_t + \alpha_2 + \psi_2 d_t)$ for the nondurable consumption, c_t , and weight $(\alpha_2 + \psi_2 d_t) / (\alpha_1 + \psi_1 d_t + \alpha_2 + \psi_2 d_t)$ for the housing services.

C.2 Welfare: Financial literacy and shocks—Top tercile

The effect of financial literacy on the welfare of the top tercile should not be large. As explained in Section 4.3, those at the top can overcome low financial literacy with their resources, although not in shorter periods, such as a few years after the shock, as shown in Figure 14. For the parameters and grid used to solve the model, the welfare loss for financially illiterate parents when facing a fertility shock in the top tercile is 5.9%. At the same time, for the financially literate, it is 8.3%. These magnitudes are comparable to Table 9 for the first and second terciles. A higher opportunity cost explains the larger welfare loss for those with higher financial literacy.

Another parameter that explains the higher welfare loss for financially literate parents at the top is the grid used to solve the model. Using a finer grid for the financial asset (b),

Table 10: Welfare loss by a shock (finer grid)

Group	Low Wealth	Med. Wealth	High Wealth
Low fin. lit.	-0.076	-0.090	-0.060
High fin. lit.	-0.028	-0.015	-0.052

Planned vs unplanned births by wealth and financial literacy level. Comparison in nondurable consumption units, equation (19). Grid for b with 201 points.

101 versus 201 points, the direction of the welfare losses stands. Table 10 shows the results for the analogous calculation of Table 9, but with a finer grid for b . From these results, financial literacy helps households to diminish the welfare losses of fertility shocks: 63% ($= 0.028/0.076 - 1$) for the bottom of the wealth distribution, 83% ($= 0.015/0.090 - 1$) in the middle, and 13% at the top ($= 0.052/0.060 - 1$). These results align with those from Table 9, in which the largest “protection” occurs in the middle, followed by the bottom, and the least at the top.